conference abstracts and applications

SIGGRAPH 98

celebrating 25 years of discovery

Computer Graphics Annual Conference Series, 1998
A Publication of ACM SIGGRAPH

conference abstracts and applications

Edited by:

Scott Grisson
Chair
Educators Program

Janet McAndless
Chair
Enhanced Realities

Omar Ahmad
Chair
Digital Pavillions

Christopher Stapleton
Co-Chair
sigKIDS/
Community Outreach

Adele Newton
Co-Chair
sigKIDS/
Community Outreach

Celia Pearce
Chair
Panles

Ryan Ulyate
Chair
Interactive Dance Club

Rick Parent
Chair
Sketches

SIGGRAPH 98

celebrating 25 years of discovery

Computer Graphics Annual Conference Series, 1998
A Publication of ACM SIGGRAPH

Conference Abstracts and Applications

COMPUTER GRAPHICS
Annual Conference Series, 1998

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

ISBN 1-58113-046-5
ISSN 1098-6138
ACM Order No. 435983

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 exhange 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 the 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.

Contents

Educators Program

Introduction

The Educators Program is a collaborative community of teachers who share an enthusiasm and love for education. We use computer graphics to teach in all disciplines, at all levels (K-12 through graduate school) and in many countries. We exchange ideas and provide encouragement to each other by sharing our successes and challenges.

Educators Program Committee

Jodi Giroux
The Allen-Stevenson School

Mk Haley
Walt Disney Imagineering

Chris Carey
Orange County Public Schools

For SIGGRAPH 98, the Educators Program features an invited talk by Randy Pausch during the opening session: "Using Computer Graphics to Unleash Creativity in the Classroom." Randy Pausch is an Associate Professor of Computer Science and Human-Computer Interaction at Carnegie Mellon University. He was a National Science Foundation Presidential Young Investigator and a Lilly Foundation Teaching Fellow. In 1995, he spent a sabbatical with the Walt Disney Imagineering Virtual Reality Studio, and he currently consults with Imagineering on interactive theme park attractions.

The Electronic Schoolhouse is a new program this year. Authors schedule times at a computer, and attendees stop by for informal discussions and demonstrations of innovative software and techniques. Software demonstrations include a HyperCard stack to teach history in middle school, an interactive Web site at the Science Museum of Minnesota, an online course on desktop publishing, and several interactive CD-ROMs for teaching science.

We have an exciting collection of panel discussions with experts in their fields. Topics include collaborations between industry and education, community colleges collaborating to teach multimedia, and a novel approach to using technology in the K-12 classroom. A workshop is also offered: "Concept Development for Computer Animators."

Papers include techniques and experiences teaching science, art, history, computer graphics, and multimedia. A unifying theme among several of the papers is collaboration between students, between disciplines, and using distance education to collaborate across continents.

Scott Grissom

Scott Grissom
Chair, Educators Program
University of Illinois at Springfield

Desktop Publishing: An Online Distance Learning Course

Elizabeth A. Hornak
Rochester Institute of Technology
jurkowski@cis.rit.edu

This online educational course in Desktop Publishing: The Basics of Design delivers course content to distance learners as a partial fulfillment of the MFA in Computer Graphics Design at Rochester Institute of Technology. The course is divided into three sections: Elements of Graphic Design, Principles of Graphic Design, and Typography. Elements of Graphic Design covers the use of color, line, shape, size, space, texture, and value. Principles of Graphic Design covers balance, contrast, emphasis, rhythm, and unity. The third section covers serif, sans-serif, display, and script typefaces, and parts of a letterform. Each section contains lecture material, as well as three printed examples and a small animation to visually illustrate the topic.

The course is designed to be completed in a 10-week academic quarter. Although much of remote learning is self-paced, a syllabus was created to provide project due dates. There are no tests, but there are four assignments. Three of the four assignments are developmental exercises arranged to correspond with the lecture material and are to be completed in two weeks. The fourth assignment is a self-initiated project that allows students to explore their own design ideas. As a companion to the assignments, there are six corresponding design examples.

Projects can be turned in for grading via postal service, the Web, or a dedicated server where files can be uploaded by the student and downloaded by the instructor. Student-to-student and student-to-instructor interaction is important and could be handled in a variety of ways, from email or whiteboards to digital audio or video programs.

The class is designed to be software- and hardware-independent. A lot of graphic design is completed with programs such as QuarkXPress and Adobe Pagemaker. Because these software packages are very expensive, it would be unfair to require students to purchase them to complete their assignments. Microsoft Word and ClarisWorks allow students to integrate text and graphics fairly well, so students are not required to purchase additional software.

The Web site was designed with sensitivity to the different computer platforms and internet browsers. Color choices were particularly important, because of different operating systems, browsers, and computer platforms display color in different ways. Colors were chosen based on Coloring Web Graphics, by Lynda Weinman.

The site was created with third-generation Web-site design techniques developed by David Siegel. Third-generation site design uses typographic and design techniques that allow the designer to control the placement of the visual elements. Invisible tables were used to give the appearance of a two-column grid and permit better integration of type and image. Another technique incorporated the use of a white or black single-pixel gif files to create indentations and align text and graphics.

Based on student evaluation, this course has a lot of utility as a distance learning class. Administering a class on the Web mitigates the problem of time and distance for many students.

University and Industry Partnerships: Creating Multimedia Solutions to Solve Unique Industry Problems

James L. Mohler
Purdue University
jlmohler@tech.purdue.edu

The goal of communication is for a message to be conveyed, comprehended, and either applied or acted upon. The ramifications of this sequence include not only the communication process (information – sender – channel – receiver), but also retention and application of information in various situations, which signifies that a transfer of knowledge has successfully occurred. Information presentation without cognitive incorporation, comprehension, or application is meaningless.

Planned interactions are known to have a very positive effect on learning and retention. Learning theorists proclaim that to reach an objective it must be practiced to help the learner cognitively incorporate it. The interaction, or "doing the objective," helps the learner reach the objective and recall the information, skill, or behavior that was practiced. Similarly, retention increases when receivers of information are able to interact with the information, particularly when the information is presented visually. Interactive multimedia requires internal user processing and focuses on the needs of the user, thereby requiring the user to actively think about the information being presented, make predetermined decisions, and presumably acquire the information or skills being presented.

Information becomes powerful, becomes knowledge, when it gains personal interpretation, comprehension, meaning, retention, and use. The purpose of interactive multimedia is to personally transfer a meaning or message from one individual to another. Traditional information distribution has been insufficient in this area – transferring the real message behind the abstract letters and data through interaction. It is through interaction that information is internalized and becomes knowledge that is alive and useful within the individual. This makes interactive multimedia a powerful medium for education and training. It is also a very adaptive tool in marketing situations, where persuasive flair helps change an attitude or belief.

Educators and presenters alike have known for a long time that information is much more readily comprehended and assimilated when it is tailored to the audience. In traditional media, a linear progression from simple to complex normally used to accomplish this purpose. However, traditional communication media cannot be everything to everyone. Writers and educators alike must cognitively organize and structure information according to their own constructs, leading the reader from what they deem "simple" to what they deem "complex." Therefore, books or other devices may be too difficult or too base for certain individuals.

Interactive multimedia, on the other hand, provides an avenue for creation of materials that may reduce the author's need to assume certain characteristics of the audience. Multiple levels of depth can be provided to satisfy various skill levels of students. If the student has no background knowledge, the lowest level of entry may be used as the starting point. If some requisite knowledge already exists, the student may begin at an intermediate or advanced level as appropriate. Interactive multimedia improves the possibility of matching the needs of the user with the available content, allowing the communication tool to be several things to several people.

Using this paradigm, over the past two years the Department of Technical Graphics at Purdue University has been developing interactive materials for several national and international companies with relative success. Products have ranged from educational and training products to marketing CD-ROMs. Many companies are requesting the department's assistance in developing multimedia- and hypermedia-based solutions that emphasize interaction.

Basic Fluid Power

Bethlehem Steel Corporation's 80-inch hot steel rolling mill training facility wanted to integrate interactive media into its fluid power training materials for employees, to improve training efficiency, effectiveness, and impact. As more and more employees quickly approached retirement, materials were needed to pass the existing employee knowledge base from one generation to the next, as well as to document newly upgraded equipment in the Burns Harbor plant. Interactive multimedia materials on CD-ROM were identified as the means to ensure the transfer of this knowledge.

CD-ROM development extended from April 1995 to December 1996. The immediate need was centered on educational materials that describe the functioning and maintenance of newly installed fluid power controls and equipment. The first CD-ROM focused on basic fluid power theory, to present the underlying conceptual information.

Content developed by the School of Technology's Mechanical Engineering Technology department was combined with the multimedia skills of the Technical Graphics Department to create the final CD-ROM.

The interactive multimedia training project provides Bethlehem Steel employees with multimedia-based informational materials that include text, sound, animation, and video to efficiently and effectively train employees in applied fluid power. Through the CD-ROM's graphical user interface, users can easily access the information and graphics, animation, and video that document and describe the system. Because interactivity is key to maintaining interest, motivation, and understanding, several of the components include interactive exercises.

SOP Interactive

As a direct result of earlier work, the Department of Technical Graphics was contracted to create interactive multimedia training materials for another plant within the Bethlehem Steel Corporation Burns Harbor facility. This project focused on the plant's Standard Operating Procedures (SOPs) and Job Work Instructions (JWIs). The existing training system was composed of text-based, step-by-step instructions that described specific processes for mill operation and functioning, as well as the order in which tasks were to be performed. Using the existing system's content, the new training materials included both text-based instruction and graphical representations of the text. Using static and dynamic graphics, the new training system provides instruction that is accessible to both readers and non-readers.

The project began with a prototype of the system, which included three SOP's. Text from the existing database system was extracted and storyboarded, and technologies for visual representation were selected: three-dimensional animations, digital video clips, virtual reality video clips, and static images.

Black Cultural Center Virtual Visit

In addition to training and educational materials for external corporations, the department has also created several interactive programs for various departments and organizations within Purdue University.

As a resource for African-American students on the Purdue campus, the Black Cultural Center (BCC) was originally designed to provide students with functional and meaningful linkages with the academic components of Purdue's comprehensive educational programs. The original BCC building was a residential dwelling built in 1905. Over time, the center outgrew that building's confines, and plans for a new building were created by an external architectural firm. To help secure funds for construction, BCC representatives needed a way to show potential donors the new building. The department of Technical Graphics was asked to produce an interactive program that could be used to market the new building before it was built.

Using three-dimensional animations, digital video, and interactive multimedia technology, the department created a presentation that included a unique graphical user interface to present both the cultural aspects of the project and the majestic architectural structure that would soon house the center. The program successfully links digital video clips of the center's current activities with digitally rendered scenes of the new building, to provide a glimpse of what the new building would look like and how it would function. The program also includes dynamic media elements that show nine of the various rooms, including a reception area, art gallery, ornate balcony, and auditorium. Current feedback indicates that the product has been the biggest single factor in the amount and number of contributions.

Conclusions

Through various industry-university partnerships, many students in the Technical Graphics Department have received real-world experience in developing interactive multimedia materials and media assets. Undertaking externally and internally funded projects provides innumerable advantages to students. It solidifies the education they receive through application, regardless of the institution they are attending. Faculty involvement in such projects increases faculty experience and real-world concepts that are conveyed in the classroom and has a direct effect on the quality of education and students.

Moving Mountains: Using Interactive Graphics to Teach Geography

Joshua E. Seaver
Special Projects Web Developer
Learning Technologies

This presentation explores the Science Museum of Minnesota's approach to informal science learning by focusing on how its classroom students and visitors to one of its Web sites, The Greatest Places Online, are developing interactive simulations to learn geography and sciences related to geography. The Greatest Places Online is an educational resource that complements the Science Museum's latest large-format film on geography for Omni-theaters, entitled The Greatest Places. The film, which debuted in February of 1998, features seven diverse geographical locations: the Amazon Basin, Greenland, Iguazu Falls in Argentina, Madagascar, the Okavango Delta in Botswana, and the Chang Tang plateau of Tibet.

In Moving Mountains, attendees have the opportunity to view several examples of multimedia projects created by students and educators from the museum for The Greatest Places Online. Many of these projects are based on construction kits whose geographical settings come from the seven places featured in the Greatest Places film. They are created using the award-winning and innovative authoring software for kids, MicroWorlds (from LCSI), where graphics, video, and sound can be brought together interactively through a simplified version of the LOGO programming language. Completed projects can be viewed through a free player that can be downloaded from the Greatest Places Web site; there is also a beta version plug-in available to view games directly from the Internet.

We believe the interactive graphics approach demonstrated through The Greatest Places Online could prove effective in classroom and home settings as well. For example, to help students understand weather systems, a geography teacher might want to teach iceberg formation as a tie-in to the Greatest Places film. After visiting the Greatest Places Online, this teacher, whose institution already owns MicroWorlds, might suggest class interactivity to help in understanding how an iceberg is formed. Students could download a basic MicroWorlds kit from the Web site, which includes a background and several "costumes" that can be attached to moveable turtles (sprites) on the screen. A variable slider that changes temperature or snowfall directly affects whether an iceberg separates from the larger ice mass. Once the iceberg separates, the movement is animated. Several icebergs can be calved and moving on the screen at the same time. When they hit each other, they bounce away. Time might also be a variable, where the iceberg slowly melts. This process can also be animated.

The educational benefit of this experience is broad. Students learn more about a geographical process. They use logic to create the variables and artistic ability to create the graphical elements. And because of the time it takes to create the project, they are more likely to let the concepts sink in and share their understanding with others.

There are plans to have an educator from the museum travel to schools and museums to assist educators in using these resources. A collaboration is already underway between the Science Museum of Minnesota and the Carnegie Science Center in Pittsburgh.

The Learning Technologies Center (formerly the Youth Computer Center) at the Science Museum has been teaching informal science classes incorporating interactive graphics since 1994. For the past four years, children and teenagers aged eight through 18 have been developing their own interactive programs as a way to learn concepts and skills in both science and art in our computer classrooms. In this process, students are asked to assemble flowcharts, storyboards, and sketches to prepare for coordination of graphics, sound, and programming. Students are trusted to come up with their own plots and puzzles, with themes given only for inspiration.

We have found that students are excited to come up with an original idea and map out how to complete it. We have also found that immersion in programming gets them beyond the desire to passively surf the Internet and engrosses them in a constructive computer activity that helps them understand logic, as well as design composition, in a fun environment. They are often enthusiastic about offering their own interactive games and animations for public viewing and for distribution to friends and family.

This process is a rich intellectual challenge and whole-brain experience, as students are stimulated to use both logic and creativity. Students come away with both a sense of accomplishment and a concrete example of their individual approach to problem solving. Encouragement to develop creative solutions that do not include violence, unless it is an expression of personal experience, has given birth to some insight and innovation in students as well.

Finally, the finished projects, being digital, have been copied and distributed via a Web plug-in, or via a player that does not require owning the program. This allows students to express their experiences and share their knowledge, which often prompts them to devote more time and attention to their projects.

Science Museum of Minnesota
30 East 10th Street
St. Paul, Minnesota 55101 USA

+1.612.221.2523
+1.612.221-.528 fax
joshuas@smm.org

We believe our approach has been successful for these reasons: students are eager to return to class, often reluctant to leave at the end of day; parents report that students are talking about their projects enthusiastically at home; parents ask how to purchase the software used. MicroWorlds, used in this demonstration, is perhaps the most widely requested software in our four years of classroom teaching.

Visitors to this presentation are welcome to discuss our museum's experiences with such topics as helping students to become self-motivating and self-organizing, techniques in programming logic, and final distribution alternatives. Participants are also welcome to share their own experiences with programming and graphics in an educational setting and discuss what the future may hold for interactivity in their classrooms.

Participants will hopefully come away with a greater understanding of, or renewed enthusiasm for, the possibilities and rewards of empowering students with the vocabulary of interactive graphics, as well as some additional knowledge of how to make this a reality in their own classrooms and homes.

The Greatest Places Online:
www.greatestplaces.org

The Science Museum of Minnesota:
www.smm.org

Science Museum of Minnesota's Learning Technologies Center:
www.smm.org/lt

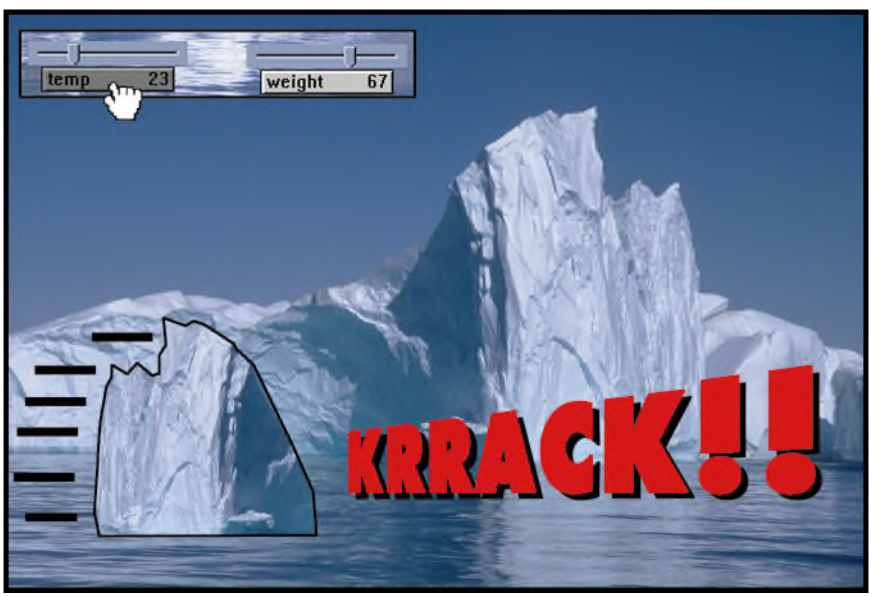

Zoom Into the Past: Illustrating History in Middle School

Karen Kresge
Argyle Middle School
Montgomery County Public Schools
Rockville, Maryland USA
karen_kresge@fc.mcps.k12.md.us

This is an adaptable, digital, interdisciplinary, culminating activity to the study of an historical period, in which students create a short, multimedia production that appears to zoom into an architectural structure of that era to a close-up of an object inside.

Sample Lesson Plan
Create A Colonial Architecture HyperCard Stack
Suggested grade level: 8

Objectives:
National Art Education Standards of Art: Students understand and apply media techniques and processes. Students understand the visual arts as a basic aspect of history.

Maryland Art Education Goal: Students develop an ability to organize knowledge and ideas for expression in the production of art. Students reveal an awareness of structures and decor made of regional material, customs, and inventions in colonial America.

MCPS Social Studies Objective: Students trace the evolution of the colonies up to 1763.

Materials: HyperCard 1.0, 29 Mac LCIII's, a VCR, TV monitor, cables, LTVPro video out card, a Laserwriter Pro printer, HD floppy disks, video tape, and teacher-created worksheets for storyboards, sharing records, and self evaluations.

Pre-assessment: Students had learned Hypercard technology and participated in field trips to colonial Annapolis and St. Mary's, Maryland.

Opening Set: [Day 1]
Students analyzed and discussed commercial slides of Colonial Williamsburg's exterior and interior architecture, tools, and decor of: the candle makers', gunsmith shop, musical instrument shop, cabinetmakers', printing and post office, bookbinder, basket weaver, windmill, apothecary, millinery, wheelwright, cooper shop, music teacher's shop, spinning and weaving, peruke maker, and Raleigh Tavern Bakery.

Project Design: [Day 2]
The instructor modeled creation of a sample, zoom in HyperCard stack and displayed the storyboard for it. Students sketched storyboards for their own colonial stacks.

Production: [Days 3-5]
Students created: on card 1, a name or title; on card 2, the exterior of a colonial public building; on card 3, an interior room of this public building; and on card 4, an extreme close-up of one object visible in card 3. Students created and placed navigational buttons that move "to next card" with a chosen visual effect on appropriate cards. Optional: students created more cards and directional buttons as motivated. Students printed storyboards. Students recorded stacks on VHS tape. Other options that require more time, which we did not do: Students record and place audio buttons on cards 2-4. Students record colonial music on a cassette tape and edit the music tape to the video tape. Students compose original music characteristic of the colonial period on a keyboard and record on a cassette tape to edit to the video tape.

Synthesis and Conclusion: [Day 6]
Students viewed and shared completed stacks in pairs or teams. During sharing, viewers related and recorded what they liked best about each stack.

Assessment and Evaluation: [Day 7]
Student recorded self-evaluations of their own projects, including: a description of the stack, the process of creating the stack, an expression of feelings about the finished stack, speculation as to what other things could have been included or improved, given more time. Storyboard sketches, stacks, and records made during sharing and self-evaluations may be used to determine the success of the lesson and plan future lessons relating to this one.

Alternatives to students creating Hypercard stacks include creating HyperStudio stacks; slide presentations in Kid Pix, Claris Impact, Microsoft Power Point, or Astound; or interactive Web pages in HTML, Claris Home Page, Adobe Page Mill, or Microsoft Front Page. Alternative images to students creating original drawings: import digital photos, scanned still prints, or digitized video taken of exteriors, interiors, and objects while on the field trip; import copied Web images, scanned text images, or copied electronic images from CD's with appropriate footnote and bibliographic citations.

Sample illustration by eighth grade student Erin Hallen.

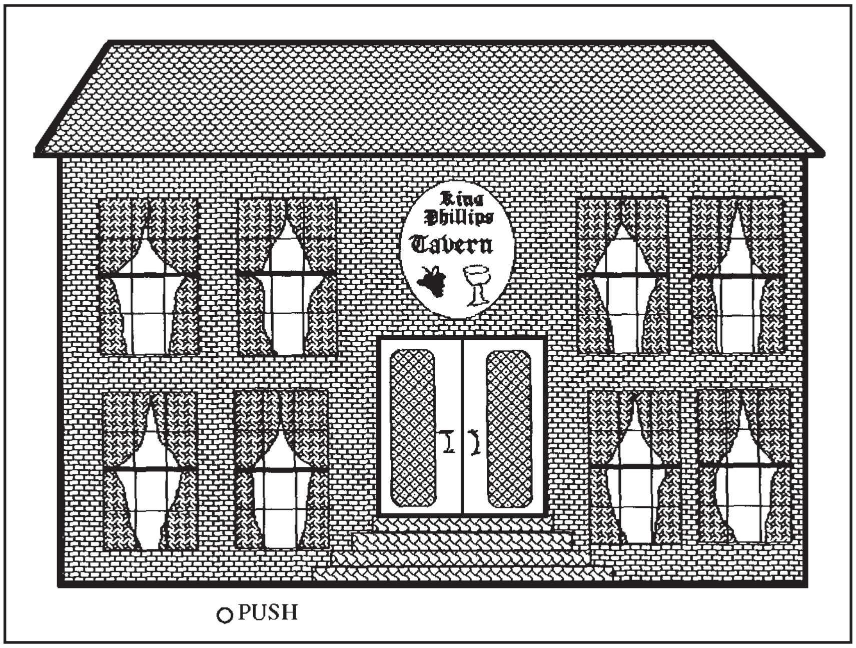

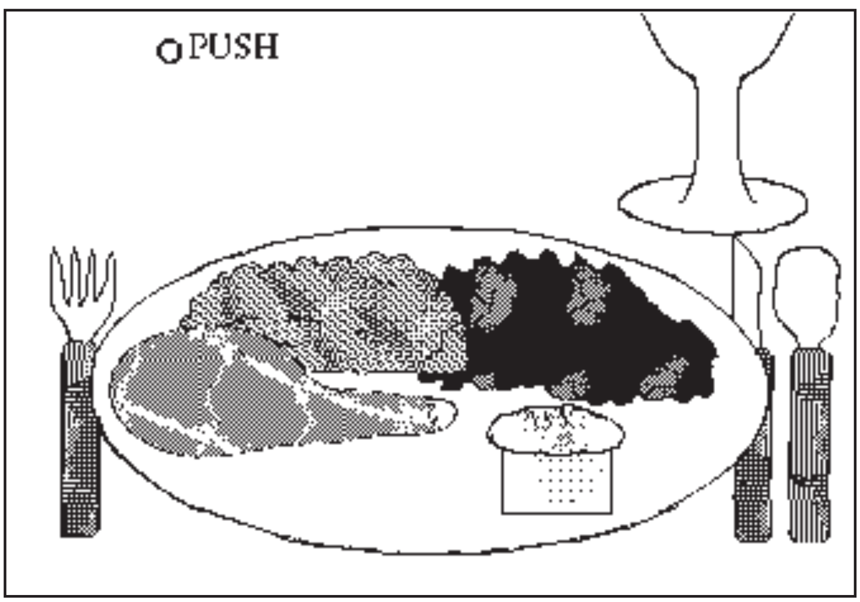

Collaborations in Higher Education

Organizer
Andrea Polli, Chair
Robert Morris College

How can art and design institutions bridge the gap between the academic environment and the "real world?" How can these institutions have an active role in the community at large? How can technology be used to facilitate collaborative projects in higher education? What are the promises and perils of integrating art and technology for collective action?

Innovative educators from several disciplines, including media and sound arts, architecture, and performance/installation discuss successful collaborations between students and cultural institutions, industry, other schools, and professional artists. Among the arts and education projects discussed: a student-designed virtual tour of a museum, a design collaboration that is taking place on the Web and in a city subway system, a collaborative-architecture design studio, and collaborations with Sun Microsystems, NASA, and internationally renowned artists including Robert Ashley, Jenny Holzer, and Guillermo Gomez-Pena.

The World Wide Web as a Medium for Improving Student Collaborative Learning
Thomas Fowler

Thomas Fowler has served as director of minority educational affairs at Cornell University and worked professionally as a designer for Davis Brody and Associates Architects in New York City and Hartman & Cox Architects in Washington, DC. He has authored several papers, including: "The World Wide Web as New Models for Capturing the Design Process," "Collaborative Learning Using Technology: Issues and Approaches," and "The Role of African-Americans in the Social & Built Environment." He has served as chair for two senior project committees and has served on 15 masters committees. In addition, he is a member of the American Institute of Architects, a design reviewer for Web Course in a Box (a software application), and an active participant in Educom's National Learning Instructional Initiative and Partners for Distributed Learning (Member of Standards for Indexing And Searching Large Digital Data Set Committee). He has collaborated with Sun Microsystems and NASA

A misconception of the Web is that it is a medium (Turkle, "Life On The Screen") that has revolutionized teaching and is most effective when used as a communication mechanism between the "learner and the teacher" (where the teacher and learner are in separate locations). Revolutionary aspects of the Web have more to do with changing the way students learn (as opposed to teaching). There are also pedagogical issues that should be addressed regarding the vision of how content should be configured for the Web environments to provide the most stimulating frame for student collaborative learning.

Regardless of the teacher-to-learner relationship, well-conceived content for the Web medium should:

- Capture the process of learning.
- Provide an inclusive environment for all students' voices to be heard.
- Create an environment for continuous dialogue.
- Establish a learning environment that evolves based on the "living and breathing" information from learners
- Develop content that is structured to adhere to medium framework rules of thumb: fast (pages between 30-50 kb), fluid (clear interface), friendly (graphic balance), fun (some relief from content).

Two diverse examples of Web content configuration are:

1. Development of a Collaborative Interactive Design Studio (CIDS, suntzu.larc.calpoly.edu/cids/) to supplement the in-person classroom student learning for a third-year architecture design studio.
2. Development of a Web distance-education course, where the Web medium is the only linkage between "teachers to learners" and "learners to learners."

The Web is exposed to the students as a medium to use in the design process. The Web pages become the meeting place where students, and others from outside the classroom collaborate and provide reactions to work in progress. Web page templates are provided to students to allow for easy configuration of design projects. Web tools provided help students capture project critiques.

Panelists

Mary Murphy
The Chicago Cultural Center

Thomas Fowler, IV
California Polytechnic State University

Neil B. Rolnick and Branda Miller
Rensselaer Polytechnic Institute

Silicon Graphics Indy Computers provide the means to generate VRML interactive computer models (online interaction and feedback during the developmental stages), and Power PCs provide graphic manipulation software for Web page configuration. Multiple venues of discussion (on-line reading analysis, critiques, weekly journals), provide a wonderful stimulus to classroom sessions.

standard professional collaborations are difficult. The Web allows invited and uninvited participation from outside the university in discussions and critiques. Early collaborations with NASA and Sun Microsystems are providing additional tools to make online exchanges more interactive.

Effective use of the Web to improve collaborations is not a panacea for

The ideas for design studio projects evolve out of working with different groups of students each quarter, and each quarter builds on the previous quarter's acquired archive of information. Projects often include both a Web presence and a parallel analog mode (physical models, drawings, etc.). These two medium extremes complement each other to provide students a diverse range of "other ways of seeing" the design process. Due to California Polytechnic State University, San Luis Obispo's isolation (four hours from both San Francisco and Los Angeles),

improving collaborative learning. Having a broader vision of the relationship between this medium and configured content is important to simulate learning. The Web does allow for introduction of real-world issues in a "living and breathing" manner into the academic learning environment. Students in the classroom start to understand that the world outside of the classroom is not constrained by the structure of a syllabus. And students connected via the Web understand that there are effective means of communicating that don't require in-person discussion.

Collaborations in Higher Education

Info-Myths and Dreams: Art, Technology, and Collective Action
Branda Miller and Neil Rolnick

An internationally recognized video artist, educator and media activist, Branda Miller focuses on the relationship among art, technology and community, and experiments with media to support independent voices, explore new visions, and stimulate critical perspectives. Miller's media art works have been screened at festivals, museums. and exhibitions; broadcast nationally and internationally; and used extensively in community organizing and education.

Professor Miller is recognized for her development of media literacy and community education projects using electronic arts and media production. For the past four summers, she has led workshops for the Five College Institute for Media Literacy. She has collaborated with several community and youth groups around the country in empowerment video workshops, including the award-winning "We Have the Force," "The Birth of a Candy Bar," "'talkin' 'bout droppin' out!!!," and "What's Up?" She is also co-editor of the manual TV EYE: Media Analysis and Independent Production, distributed by Boston Film Video Foundation. Her media activist projects include producing/editing "Art of the State/State of the Art?," "National Arts Emergency," and "Cori: A Struggle for Life." She developed free distribution, with accompanying curricular support, as part of the design for using media art as an educational tool for community empowerment.

Neil Rolnick has been active as a composer and performer of computer music since the late 1970s. He performs on a portable computer music system and concertizes regularly in a wide variety of contexts throughout North America and Europe. He has appeared as featured soloist with ensembles such as Dogs of Desire, The California E.A.R. Unit, Relâche, Gerard Schwarz's Music Today Ensemble, Musical Elements, Gamelan Son of Lion, and the Albany Symphony Orchestra. Current compositional projects include an evening of music for voices and amplified chamber orchestra for Dogs of Desire, which is to be recorded and released by Albany Records in late 1997; a series of new works for improvising chamber ensemble; and a series of new solo performance pieces.

Rolnick studied musical composition with Darius Milhaud at the Aspen Music School, with John Adams and Andrew Imbrie at the San Francisco Conservatory, and with Richard Felciano and Olly Wilson at the University of California at Berkeley, where he earned a PhD in musical composition in 1980. He studied computer music at Stanford University with John Chowning and James A. Moorer, and worked as a researcher at IRCAM in Paris from 1977 to 1979. At Rensselaer Polytechnic Institute, Rolnick has directed creation of a unique Master of Fine Arts program in electronic arts, which focuses on a truly integrated approach to time-based art and performance with the electronic media.

Art practice, integrally connected to cultural and social action and experience, is the creative avenue for critical awareness, independent and collective action. It lets us reflect upon the possibilities of our bodies, communities, and the space we exist in, and explore our relationship to history and time. Computer technologies have produced revolutionary change in the way we experience our bodies in space and in time, construct our identities, and connect to the world outside our horizon. With digitized fluidity, image, sound and text are woven into hybrid fantasies.

The iEAR Studios, always committed to hybrid art-making in the advanced technological landscape of Rensselaer Polytechnic Institute, has produced many collaborative projects that offer aesthetically and critically challenging models for bringing together art and technology, education and community. These projects explore new possibilities in diverse forms of electronic media, while examining the political and social manifestations of these explosive technological accomplishments.

How does the computer promise to offer more information? How is that information constructed and delivered? Who has the power not only of transmission but of real "mass" distribution? How are we educated as the receivers of that information to critically analyze and apply this information to meet our needs, not only as consumers but as empowered citizens, artists, and dreamers?

We refer to a range of works, including iEAR collaborations with internationally renowned artists including Robert Ashley, Jenny Holzer, Guillermo Gomez-Pena, Alvin Lucier, Muntadas, Tony Oursler, Pauline Oliveros and Ione; public arts events and arts and education projects located in the historic "birthplace of the Industrial Revolution" in Troy, New York; and multimedia performances and innovative experimentation with interactivity, as we share and critique the "Info-Myths and Dreams" of the Computer Age, and promises and perils of integrating art and technology for collective action.

The Great Chicago Cultural Center Adventure: An Interactive Tour of the Cultural Center

Andrea Polli and Mary Murphy

Andrea Polli is an electronic media installation and performance artist who creates public performances and installations with aural and visual components. Her live system for musical improvisation between a human performer and a computer was presented at the Mathematical Association of America/American Mathematical Society Joint Conference at the Ferienkurze für Neue Musik in Darmstadt, Germany. She has recently received a Chicago Artist's Assistance grant for her current work, "Virtual Space and the Construction of Memory," published in Leonardo, April 1998, and presented at the 1998 Imagina Conference in Monaco.

She has worked internationally in performance projects, most recently with the artists' groups SEL (Super Ex Libris) and WALK, on a large-scale installation and performance project in a 200-year-old observatory in Vilnius, Lithuania. As a member of Artemisia Gallery, she has organized several national and international exhibitions: Ada, a celebration of women interactive artists, and Meme Me: Identity in the Replication Age, a project in conjunction with ISEA97 (the International Symposium of Electronic Art). She is a paper reviewer for the SIGCHI conference and a member of the American Institute of Graphic Artists and the American Center for Design.

As chair of the Robert Morris College Institute of Art and Design, Polli has helped to develop innovative curriculum for electronic art. She has created collaborative environments for students, faculty, and artists both in and out of the lab. Her most successful collaborations include The Great Chicago Cultural Center Adventure, an interactive CD-ROM created by students of Robert Morris College with the Chicago Cultural Center, and Live live! a public art project with RMC students, The Museum of Contemporary Art, and the Chicago Transit Authority (www2.rmcil.edu/live).

Mary Murphy is the education director of the Chicago Cultural Center, one of the most diverse and innovative centers for the arts in Chicago. The center presents hundreds of free public events each year. The education department serves the entire Chicago Public School system, as well as hundreds of private schools in the Chicago area. It is well known for its cutting-edge ideas and programs. Murphy is also an active freelance arts writer and teacher who has taught at the School of the Art Institute of Chicago, and she is a board member of Randolph Street Gallery.

In February 1997, the Robert Morris College Design Institute and the Chicago Cultural Center Educational Department began a city-wide educational project utilizing interactive technology. The RMC Design Institute has three areas of concentration: Graphic Art, Computer Aided Drafting, and Multimedia and Computer Imaging. The focused study of the Multimedia and Computer Imaging concentration guides exploration into every aspect of the process of designing for interactive multimedia. Students participate in initial planning and research, identifying the target group(s) for the specific project, and creating storyboards and flow-charts of interactivity, which are reviewed and revised based on design requirements and aesthetic concerns. In the lab, students then create the interactive work in design teams.

After a model of the project is created, the students participate in user-testing research. The model-testing project is essential to give students a full understanding of the multimedia design process and to assure the effectiveness of the finished product. It is during the testing process that students explore design theory and practice. The Design Institute also hopes that follow-on research will help to further the study of interaction design as a new design discipline.

A team of two instructors in the Design Department selected a class of 15 RMC Multimedia and Computer Imaging Associate Degree students to participate in The Great Chicago Cultural Center Adventure, a collaborative design learning and production process. Working closely with Mary Murphy of the Cultural Center's Education Department before the project began, they defined the goals, basic content, and parameters of the project. The project quickly grew. An RMC Computer-Aided Drafting instructor brought students into the project to help with planning and publicity. A video instructor created a video documentary of the Chicago Cultural Center, and the graphic art area designed print material for the project.

Project Overview and Goal

Produce an interactive computer application that can be written to a CD-ROM, which will serve as an introduction for school children to the Chicago Cultural Center. The final goal is a CD-ROM that will increase students' enjoyment and appreciation of the history and architecture of the Chicago Cultural Center and the role it plays in the cultural life of Chicago.

Much of the information is based on the existing curriculum tool developed by the Education Department. The narrative focuses on three teenagers who sneak into the Cultural Center. Each character represents one of three main sections of the CD-ROM:

1. Knowledge, which focuses on the half of the Cultural Center that was originally built to house the Chicago Public Library. Topics include literary history, the Chicago Public Library, and the different things the space has been used for in the past 100 years, including current cultural activities. The navigational structure of this section is based on the experience of being in a library: pulling books off a shelf, turning pages, etc.

2. Conflict, which focuses on the half of the center built as the Grand Army of the Republic's Civil War Memorial. Topics include Civil War history and the changing uses of that side of the building. The navigational structure of this section is centered around two main ideas: place and personality.

3. Structure, which focuses on the architectural history and details of the entire building. Topics include architectural terms and styles found in the building, construction and renovation, and historic proposals for alternate uses of the building. The navigational structure of this section uses spatial and temporal navigation via an elevator/time machine.

Each of the three main sections contains several activities, including animated sequences and interactive games or stories. The completed prototype was tested with sixth- to 12th-grade students to evaluate the effectiveness of the content and structure of the guide.

Reading, Writing, Reload: New Three Rs for a New Millennium

Organizer
Sarah Feldman
Thirteen/WNET

Somewhere between pariah and panacea lies the fact that, when used strategically and creatively, interactive media can be an invaluable tool for instigating thinking, investigation, and creation. This panel explores the opportunities and myriad uses online and interactive media provide educators, students, and parents. Topics include: designing and using graphic media to maximize its educational potential, exploiting online interfaces to shape thinking and spur constructivist curricula, using multimedia to facilitate "sneaky learning," the importance of promoting media literacy, and encouraging girls to use technology.

Context, Hypertext, What Next?
Sarah Feldman

Much of the discussion surrounding classroom use of multimedia and the World Wide Web focuses on issues of availability. While access to resources is undeniably important, assessing their value is just as critical. Too often the question "Are we using technology and media in the classroom?" overwhelms our investigations into "HOW are we using technology and media in the classroom?"

As National Project Director for WNET in New York City, I work with teachers and media professionals across the country on the strategic use of visual media and technologies in the classroom. For nine years, our National Teacher Training Institute for Math, Science and Technology (NTTI) has provided thousands of educators with the vision and techniques needed to integrate a range of media into their curricula. I'd like to share our strategies for maximizing online and multimedia's creative and instructional possibilities.

Specific strategies include how educators can:

- Instigate media literacy activities in classrooms, teaching children how to filter and evaluate information for accuracy, priority, and relevance.

- Add depth and context to curricula. While new media are (theoretically) inherently interactive, educators, users, and authors can look at ways in which content can be better contextualized and expanded – maybe even understood – through hypertextual media.

- Shift the way students grapple with information. The inherently conjunctive role of hypertextual environments presents whole new possibilities for researching, perceiving, synthesizing, and communicating information. Where once information was organized in linear or encyclopedic formats, now information is often presented and understood by way of its connection to other information. Educators can maximize hypertext's power to bring multifarious elements together with a more comprehensive, deeper sense of the associative and interdependent world of ideas.

As we look toward more student-centric and technology-enhanced classrooms, it's important that educators (content developers) examine interactive media's unique capacity to help transform how we learn, communicate, and think.

Panelists

Henry Bar-Levav
OVEN Digital

Anthony Chapman
Thirteen/WNET

Aliza Sherman
Cybergrrl Internet Media

Design for Education
Henry Bar-Levav

Designing interactive educational media involves weighing two priorities that frequently appear to be in conflict: ease of use versus short attention spans. In other words, building a Web site that is easily navigated and understood is often overshadowed by the need to feature "fun" bells and whistles to keep children's attention on the learning tasks at hand.

Designing for education means thinking about learning – as a process – in a new way. I believe that quality information and interface design instruct users on several levels as they proceed through the product. Web-based learning programs, whatever the subject matter, can simultaneously teach organizational skills, since as students learn to navigate the product they are developing an intuitive sense of its organization.

I feel strongly that, in the "digital age," we must cultivate a design-sensitive culture as part of the emerging multidisciplinary curriculum. We all agree that students should have some understanding of the basic scientific principles of the world they inhabit, such as gravity, the composition of matter, and the water cycle. I believe they also need to have basic knowledge of design issues in the world around them, not merely in terms of aesthetics, but more significantly in terms of "design" in a deeper, structural sense, as a system of organization. Interactive learning programs, when designed with these factors in mind, can be a primary vehicle for conveying this important general knowledge.

Bridging the Gap Between Games and Education
Anthony Chapman

- How many more animated films can be recycled into educational titles for kids?

- Making dresses for Barbie using a color printer has proven to be commercially viable, but what is it teaching the next generation of women?

- What does the future hold for children who need to ask mommy and daddy for new retinas for Christmas because the VR version of DOOM has fried their old ones?

In my opinion, the state of kid's software titles, especially educational titles, is at a low point. Why not just watch TV? Where are the new products that utilize the potential of today's technology? In the past two years, I've worked with educators, children's writers, and software publishers developing and producing two projects that I hope will shed some light on the topic of "sneaky learning." I would also like to explore new ways to help software find its way into children's rooms at home and in schools.

New software products could bridge the gap between games and education. It's important that educators and developers be aware of software that is currently available and being used in schools, current statistics on the level of technology in schools, and plans for making better use of technology in the classroom.

The Importance of Gender Equity in Multimedia Instruction and Production
Aliza Sherman

When girls reach puberty, their use of and interest in technology often wanes as boys become more aggressive with computer or technical equipment. Teachers need to be better equipped to insure that girls are not left behind as these changes occur. Parents also need to become better educated about ways to encourage their daughters to use technology and computers.

Luckily, a wave of new Web sites for teenaged girls is appearing, giving girls a voice and community via the Internet. Teachers need to integrate the Internet, particularly girl-specific resources, into their curricula to help encourage girls to embrace technology. Engaging girls in technical activities as early as possible will help position computers and technology courses as more interesting and accessible. Since women comprise less than half of the student enrollment in the computer sciences, it is important to find innovative ways to encourage girls to pursue more technical curricula and careers.

P.LU.N.G.E into School

Organizer
Karl Hook
Florida State University School

Panelists
Debi Barrett-Hayes
David Godwin
Marleni Young
Florida State University School

P.L.U.N.G.E. (Practical Learning Utilizing New Gadgets in Education) is one part of a comprehensive research agenda at the K-12 developmental research school on the campus of Florida State University. The school's mission includes research, dissemination, and service, in addition to meeting the educational needs of its 1,128 students.

P.L.U.N.G.E. began as an initiative to integrate technology into the middle- and high-school curricula. The idea was to introduce middle-school students to the skills required to produce multimedia projects so they could "show what they know."[1]

In addition to senior projects, students were required to develop a portfolio of the works that best reflected what they learned. Traditional portfolios are increasingly cumbersome – difficult to manage and replicate. So the school developed electronic portfolio assessments.

When student work is saved on an electronic storage medium, portfolios are easier to manage. Objects can be added and deleted, stored and duplicated, very quickly and efficiently. When they apply for employment or admission to institutions of higher learning, students can copy their portfolios and send them as part of their application packets. Unlike "paper and pencil" applications, electronic portfolios allow students to send recorded sounds, images, and movies that provide reviewers more insight on the applicant.

The key to P.L.U.N.G.E. is not the product, but the process. As they developed strategies for showing what they know, students became more familiar with the technology. More importantly, they increased their school-to-work skills: planning, problem solving, flexible thinking, working independently and collaboratively, and critical thinking.

In order for the P.L.U.N.G.E. concept to be fully realized, technology can not be isolated. Computers and other electronic tools cannot be limited to special elective courses or used solely for remediation, enhancement, or as a reward for those who finish first. Technology must be integrated into the curriculum, not to replace, but enhance, the activities taking place in the classroom.

P.L.U.N.G.E. is based on the premise that:

- Learning should be for a purpose (practical). What is learned should be connected to a real-world situation.
- Students should be encouraged to explore and try new ideas, and to take risks and challenge themselves. Through disequilibrium, learning is achieved.

* The learner should use the technology. When they are immersed in the tech-

P.L.U.N.G.E.
www.fsus.fsu.edu/~plunge

MagLab: Alpha
k12.magnet.fsu.edu

nology, students are better able to identify its strengths and weaknesses.

- Educational institutions should be ahead of the private sector in applying state-of-the-art technology. In order to facilitate this, students should be placed in externships that introduce them to the latest techniques.

- Technology should not divert the learner from the basics.

- Learning should take place in the context of an educational environment that is innovative, substantial, and supportive.

History

During the summer of 1994, P.L.U.N.G.E. was initiated to field test techniques for implementing electronic portfolio assessments. Twenty-eight students worked in teams of four to develop multimedia projects focusing on social issues pertinent to them (censorship, teenage pregnancy, teen suicide, etc.).

In the 1994-95 academic year, all eighth grade students at the Florida State University School were enrolled in a course that combined art with computer technology. The purpose of the course was to teach the students how to use the technology so that they could "show what they know" in their science classes and later use their skills in their high school classes as the school embarked on alternative assessment projects.

The multimedia class was such a success that the students began requesting it at the high school level. Since the additional course offerings did not fit the schedule, the classes were scheduled for the summers.

Through the P.L.U.N.G.E. project, the Florida State University School has increased its use of multimedia as a tool. The students have produced several multimedia products for local and state non-profit agencies and have produced public service announcements for the local cable network.

middle-school curriculum materials on magnetism. The students assisted in production of the graphics for the publication and produced "Magnets, Medicine, and Me," the first of an eventual series of videos on magnetism

Most recently, the P.L.U.N.G.E. team has been collaborating with the National High Magnetic Field Laboratory (Magnet Lab) on a number of projects, including STAR TREE and MagLab: Alpha.

During the summer of 1996, the P.L.U.N.G.E. students worked with the Department of Education's STAR TREE (Science Teacher And Researchers Translating Research Experiences into Educational Experiences) at the Magnet Lab. The purpose of this project was to produce

called the "Many Adventures of Lloyd."

In 1997, the P.L.U.N.G.E. Team collaborated with the Magnet Lab on MagLab: Alpha. P.L.U.N.G.E. students produced all of the graphics for the "Alpha Guide" and did all of the scripting, programming, and design for the Alpha Interface, a multimedia CD-ROM. MagLab: Alpha is being distributed statewide through the Florida Department of Education's Area Centers for Educational Excellence (ACEE) and will soon be distributed nationally.

P.LU.N.G.E into School

The P.L.U.N.G.E. project has also ed to student opportunities in higher education and career preparation. Many P.L.U.N.G.E. participants have been placed immediately in highly skilled positions, even as they continue their education, fulfilling our "school-to-work" initiative.

Although the P.L.U.N.G.E. team has worked on numerous projects, this is our most recent accomplishment, and it has the greatest potential for impacting education, so it has become the focus of our attention. In addition, we are hoping P.L.U.N.G.E. will become the basis for future endeavors. For those reasons, we have adopted a pseudonym: The Alpha Team.

Alpha Team member David

Through P.L.U.N.G.E. and working on MagLab: Alpha, I have been able to learn many new things and use some of the latest graphics and computer technology. Working on MagLab: Alpha helped to further my knowledge and interest in high-end computer graphics technology. The kinds of resources available to me through the P.L.U.N.G.E. team would not be accessible through a normal classroom setting.

I was given a lot of responsibility for completing my part of MagLab: Alpha. This gave me a feeling of working on a job, rather than a school project. In this project, like on a job, the better work produced, the more responsibility given.

Through working with P.L.U.N.G.E., I have learned a lot about science, computers, and graphic design, which I might not otherwise learn or even experience.

Alpha Team member Marleni

I feel that P.L.U.N.G.E. was a very essential part of my learning experience. Due to the high standards demonstrated by the P.L.U.N.G.E. Team, we were able to collaborate with the National High Magnetic Field Laboratory. This gave us the opportunity to acquire new skills and knowledge while working on a practical project. The end result would be a real product to be used in schools by middle school students.

This "work-like" opportunity allowed us to work in a lab with a wide range of technology. It was a lot easier to learn to work and run new programs that required a lot of RAM at the Magnet Lab. If we had been in a regular technology lab, the computers would not have had enough memory to run half the programs we needed. With so many more options at our fingertips, it was easy to have an excellent product. The group could focus on what we wanted the product to do instead of dealing with hardware limitations.

Also, there wasn't just one teacher or administrator we could get help from. When we needed information, there were multiple scientists and other people who could give us a first-hand look at whatever it was we needed. We also learned to use the expertise of our peers. We found that many times the students could find the solution to the problem and help the adults.

It was exciting to know that the work that we were doing as part of a class assignment would be turned into a tool for middle school students to learn from. We all wanted to make sure the work we did was as perfect as possible, because our names would be on the final product. We all made sure to put the extra effort to ensure the highest quality product possible.

The entire experience involved a creative process. We would work on a portion of the project or try a new idea and then demonstrate it and discuss it with the team. We were always stopping to review, reflect and edit, or start over, if necessary. We critiqued each other's work and sought out other options. Also, our work depended on real-world feasibilities, such as what an average school could spend to acquire this product.

After we finished working for class credit, some of us were asked to return to work for pay. This has thus led to some of us getting jobs after the summer was over.

I feel that my influence in the group was just as important as anyone else. We were all part of the team, and we each contributed something unique to the project. That has also carried over into the new school year, and we still consider each other members of the team.

I was responsible for the Translator, a section that listed terms, and gave their definition. In some cases, we also included diagrams or videos to better explain the terms. I was then responsible for the layout of the exploration section, which gave a supplementary view to that of the Alpha Guidebook. I was also responsible for completing the section on the Alpha Team Members. In this section, I used the animal characters and explained their relationship to magnets and magnetic fields. Also, I included graphics that led to a look at the adults and students involved in creating the Alpha Interface. The last section that I was responsible for was the Careers section, which described various jobs at the National High Magnetic Field Laboratory. For this section, I had to collect sound clips, pictures, and quotes from the lab's employees.

These many different experiences helped me to see the many different windows of opportunity that can be opened with the right guidance. There were many other students who passed up the chance to be part of the project in order to stay at home all summer, or to take a class that they thought would be "easier." But for me, the fact that this was something that I had to work hard at and make sacrifices for made it more worthwhile. The fact that it was a bunch of teenagers and teachers coming together with some of the most talented people in the world was just amazing! It made education a life experience instead of a class credit.

I would take this course any day! I had some of the most rewarding experiences and wouldn't trade them for anything!

1 Dana, T., Lorsbach, A., Hook, K., & Briscoe, C. (1991). Students showing what they know: A look at alternative assessments. In Kulm, G. & Malcom, S. (Eds), Science Assessment in the Service of Reform. Washington, DC: AAAS.

Where Industry and Academia Meet: An International Perspective

Multimedia systems and digital media are becoming increasingly important throughout the world. There is also a shortage of people who are educated and trained to develop, design, and implement these systems. Today's computer interfaces are designed in a much more visual and intuitive manner, so the skills involved in creating software applications have expanded from traditional programming to include creative skills, such as graphic design and writing interactive narratives. Education of graduates who are joining the information technology industry must now address the need for these new multidisciplinary skills.

This panel discusses these problems and issues in both academia and industry in Sweden, including: the convergence of traditional computer science students with traditional artists, obtaining qualified staff, and how industry is working together with academia in solving these problems. We also look at how to introduce university graduates to industry. We examine some of the industry-education problems caused by the structure of the Swedish education system and propose some possible solutions. We hope that by airing these issues for discussion, both industry and academia can understand each other better, locally and internationally, and pursue the solutions that are presented.

One of the problems that we encounter is the need to focus on communication, not the computer: the need to provide communication that connects people, companies, schools, and clients. University and industry need to collaborate to solve these problems.

Topics for discussion:

- Universities traditionally tend to focus on the process of learning, while companies look at the end product. What is better for the student?
- Students seem to want to learn products rather than theory. Why is this good and bad?
- Current trends in the Swedish industrial sector and how both industry and academia are responding.
- What is the balance between theory and practice? What sort of practice is needed?
- How do you mix the artist and the programmer? How do you train both?
- Digital media companies are often involved in strong group projects that require highly developed communication, management, and organization. Should these skills be taught at the university level?

Figure 1 Student work at Commando Royale. (Student - Mikael Håkansson, UCGS-KP96)

Mark Ollila
Mark Ollila is actively involved with the digital media industry in Sweden through his role in the Creative Programming Program at the University-College of Gävle and his position as a SIGGRAPH representative. He has always been interested in 3D and computer graphics as a hobby, but in recent years these activities have become a full-time career.

The Creative Programming Program, which began in the autumn of 1995, is unique in its approach to selecting students and in the teaching environment it provides. There is a huge emphasis on quality and creativity, which attracts a wide range of students from different cultures. The course looks at both technical and aesthetic issues involved in today's digital media. With a strong emphasis on theory combined with industry involvement, it keeps students up to date. The student selection process is quite stringent. A jury reviews initial applications and invites candidates for interviews, which generally involve examination of the applicants' analytical and logical abilities, creative abilities, computer skills, and evidence of a creative mind.

Creativity is the most important criterion: the ability to look at a problem and solve it in different ways. For this, theory is very important. It will always going be the same, as the software and hardware change. Issues in Swedish education all boil down to not having the resources to find good quality digital media instructors. Unless the industry gets more actively involved in education of students, Sweden will find itself with thousands of digital media artists with no industry-quality skills at all. Everyone can learn to use a word processor, but that does not make everyone a novelist. The same principle should be applied to animation and digital media. Universities and industry must initiate better communication and understanding of what they need and can deliver.

Another issue in education is the government's approach to funding. Education officials decide that a field such as science is worth a certain investment per student without examining the real differences in educational requirements between chemistry, physics, and multimedia. Also, they don't consider the fact that the

Organizer
Mark Ollila
University-College of Gävle

Panelists
Joakim Kempff
Commando Royale

Johan Ljungman
Silverfish

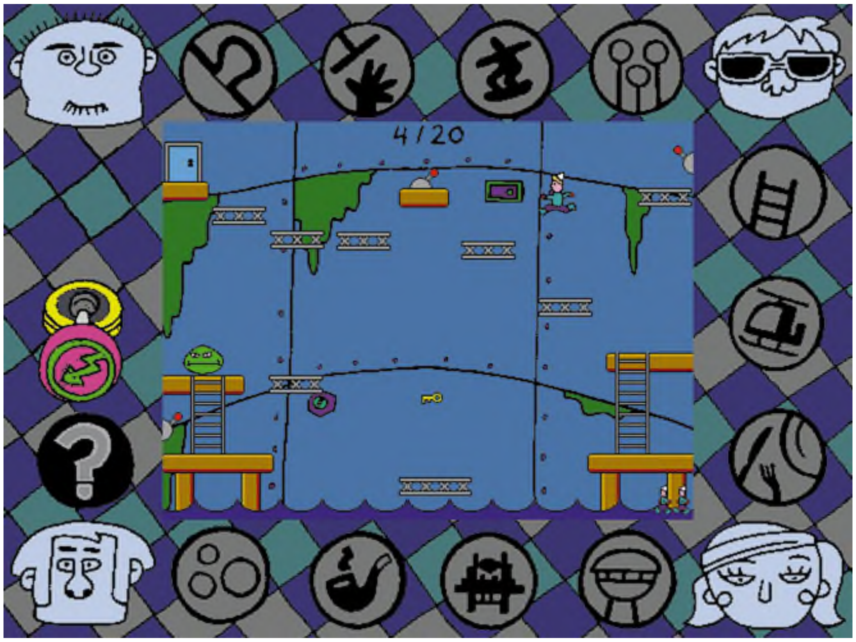

Figure 2 Student work at Silverfish. (Student - Niklas Konstenius, UCGS-KP95)

multimedia industry is growing 400 percent per year, so we need courses to provide the jobs. What about the quality? Imagine if this principle was applied to medicine: "We need doctors. You, you, you, and you shall now teach medicine." When does a computer science teacher become a computer animation teacher?

How much is one quality student worth to industry compared to 10 average students? It is time to get industry off its pedestal, and get it more involved with the education system by sponsoring students and schools, participating in courses, and offering internships, rather than criticizing universities for delivering below-par students. On the other side, the government and the universities must decide which level of "quality" they are addressing and what sort of end result are they really looking for. Is it a question of education for the masses or education for those who want to be employed? Are we educating people for jobs, or educating them in the ability to think and learn?

Joakim Kempff

Joakim Kempff is managing director of Commando Royale, a new special effects company for TV and film. Commando Royale is the largest 3D modeling and animation house in Scandinavia.

The present problem for digital media companies in Sweden is that there is going to be huge market growth in the coming years but not enough skilled people to deal with it. The big problem is to find creative people who can deal with problem solving (illustrators communicating with programmers, for example).

This indicates that the different cultures need new structures. Who is to educate the student: the university or the company? Today, companies and universities are two different cultures, with different structures and hierarchies. In the future, they have to be more the same. One big problem is that students want to learn the software rather then focusing on problem solving. As the software changes rapidly, the student is outdated. For example, a person who knew Alias three years ago would have to know the whole film industry today to be able to compete for the same job.

Problems also exist between the two cultures in research. Universities, because they have the time and money to experiment, can come up with solutions that generally will not be thought of at companies. Companies, on the other hand, have to be more efficient with their resources and may not be as adventurous in research. But the end result of the university research is ideally something that a company can use. How can we assure this end result?

Johan Ljungman

Johan Ljungman is the CEO of Silverfish Interactive Entertainment, a Projector New Media company based in Stockholm that conceptualizes and develops interactive entertainment, digital TV, online and CD-ROM games, and events.

I hear from art school students every day, asking to obtain work experience, telling me that advertising agencies have told them that they will not be let into the business unless they are already in the business. This causes inbreeding. How does the industry change? In Sweden, this business is like the pre-season players market, where every one is buying and selling players, and I am here waiting for the referee to blow the whistle so the game can start.

In Sweden, people are not usually employed based on their CVs, but on their creativity and social capacity. The fire and the hunger within the candi-

Figure 3 Student work at Commando Royale. (Student - Mikael Håkansson, UCGS-KP96)

date are of major importance, and that's where universities and industry often differ. The university stresses the CV. Industry looks at experience. I also want to point out that I don't want to be the boss. I want a team, where everyone can think, take initiative. I don't want to be the suit. I want creativity and the knowledge that anyone in the team is capable of leading the company. But that leads us to the question: what is more important, the programmer or the artist? I want a creative programmer. A year ago I wanted a programming artist. Which one is better?

Today, the government is funding hundreds and hundreds of poorly prepared, understaffed courses in multimedia. How is the school going to pay for quality teachers when industry is struggling to find qualified staff? The government is also fooling kids by saying: "Learn Director and Lingo, and you will get a dream job." This is particularly true in the eight-week multimedia courses around the country. They may be good for project managers but not for producers. Graduates of these schools have average all-round skills, and are good at communicating with the people who have advanced skills. But for every 10 staff members, we only need one all-rounder.

The situation reminds me of the mid-80s desktop publishing boom, when new companies and university courses popped up. Now, people realize that it is not the tools, but the skills. At Silverfish, we only want people from the three or four quality institutions. The other 100 universities are generally producing students of lower quality. Also, "multimedia" and "information technology" are stupid words. They're too broad. What is a multimedia degree? A multimedia degree from one university is totally different from another university's multimedia degree.

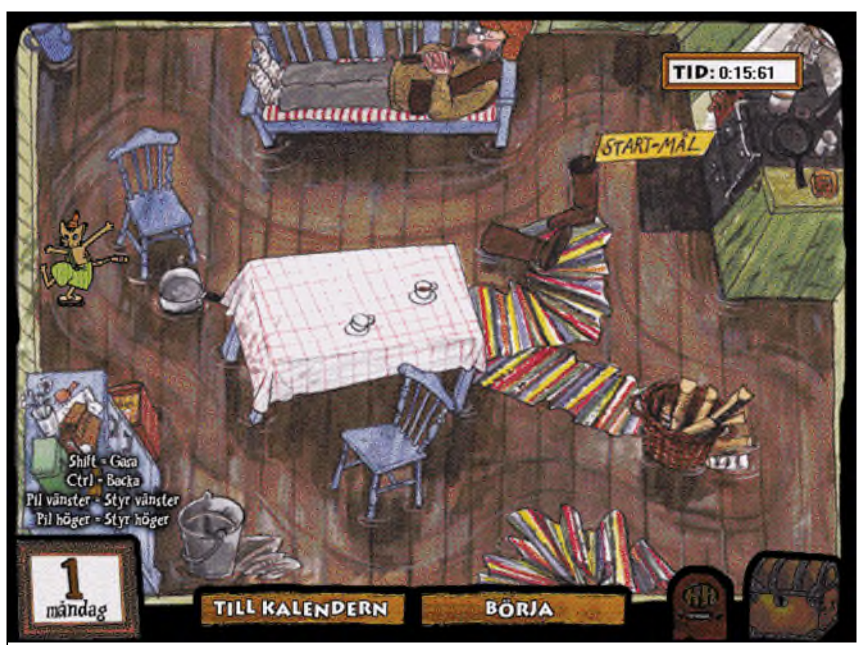

Figure 4 Student work at Silverfish. (Student - Johan Sundblad, UCGS-KP96)

Multimedia Boot Camp: Adventures for the New Millennium

Laurie Burruss
Pasadena City College

Karen Owen
San Diego Community College

Multimedia Boot Camp: Adventures for the New Millennium responds to the severe shortage of skilled workers for the nation's multimedia/entertainment industries. It encourages faculty to create and to collaborate on multimedia projects and curriculum that is pertinent to their classrooms and their areas of expertise.

When a consortium of four California community colleges, Creative Technologies Institute (CTI), expanded its successful "Accelerated Multimedia Curriculum" development effort (funded by a U.S. Department of Labor Demonstration Grant) by implementing statewide and regional curriculum and staff development in multimedia technologies, it proposed implementation of a hands-on and online experience in which faculty with expertise in diverse disciplines develop instructional multimedia curriculum that enhances student learning and creates multimedia experiences in the classroom. This project uses alternative delivery methods such as online courseware, distance learning via video conferencing, and the Internet as an instructional tool. It focuses on staff and curriculum development using multimedia tools in traditional classrooms as well as electronic delivery of multimedia instruction to students. Since multimedia lends itself to interdisciplinary project-based activities, the focus is on project outcomes and collaboration of faculty throughout California.

Implementation of Multimedia Boot Camp
In late 1997, CTI solicited brief applications from California Community Colleges interested in participating in the Multimedia Boot Camp demonstration/"beta" project. In the first demonstration of the project, in the spring of 1998, a total of 35 faculty (five from each of the seven participating colleges) comprised the first "virtual" class. Five cross-disciplinary applicants were recruited from each of the participating colleges to participate in the training program. One of the five faculty at each college was designated the lead faculty coordinator to help with information dissemination and project coordination. These participants made up the "virtual class."

The demonstration project consisted of a two-day hands-on multimedia workshop entitled Multimedia Tool Kit followed by an eight-week online course with four "adventures."

The Two-Day Hands-on Workshop
The hands-on workshops at two sites in Southern and Northern California by regional program introduced faculty participants to the Multimedia Tool Kit hardware, software, scanners, digital cameras, and A/V equipment that they would need to develop instructional multimedia. After the workshop, the faculty coordinators from each of the seven colleges introduced the faculty participants to their own school's facilities (instructional and digital labs) for support of this course.

Each faculty participant left the workshop with a Web site (minimum of two pages) designed and posted to a Web server as well as an interactive electronic presentation. Adobe Acrobat and Pagemill allowed the faculty to create Internet-ready content using familiar and available applications such as word processing, MS Powerpoint, and desktop publishing applications. In Acrobat, they were able to add multimedia elements such as audio, QuickTime, video, QuickTime VR, animations, and branching links to internal text, graphic images and charts, URL sites, and email addresses. They learned how to print a PDF file, embed a PDF file in an HTML document, project an interactive electronic presentation created in Acrobat, and download their interactive multimedia presentations in PDF format to their classrooms, auditoriums, libraries, and students.

The Online Adventures
After the hands-on workshop, four online adventures were scheduled every two weeks beginning the first week in March. The four adventures focus on the following topics:

Adventure 1	Searches on the Internet
Adventure 2	Good Design
Adventure 3	Navigation: A Trek Through Content
Adventure 4	Putting Multimedia to Work in the Classroom: The Final Ascent

In Adventure 1, participants are oriented to Top Class, the online instructional delivery system and introduced to how to participate in various aspects of an online course. The objectives of the online adventures are:

- To allow faculty from a variety of disciplines to create multimedia projects and content that is pertinent to their classrooms and areas of expertise.
- To develop multimedia instructional strategies and skills in multimedia tools and applications.
- To not only experience alternative interactive multimedia courseware via the Internet, but also to help each instructor develop four interactive multimedia lessons that can be demonstrated with their students.
- To focus on curriculum activities that routinely integrate interactive multimedia presentations in traditional classrooms.
- To enhance and support instruction with dynamic Web sites created by faculty.

Concept Development for Computer Animators Workshop

Concept Development for Computer Animators is a workshop for educators who teach time-based media using computers. One of the major problems facing young time-based media students is understanding how to prioritize concept and idea over technology and technique – how to allow technology and technique to support the idea. Without a good idea there is no piece. The workshop addresses issues specific to concept development: how do you take students through a process that will empower them to develop robust concepts that make sense to do using the computer? Participants are introduced to the analogies between computer animation and theater, computer animation and film, and computer animation and sound. Participants are taken through motivational exercises suitable for an animation class and shown examples of work from the concept development class taught at Ringling School of Art and Design.

Teaching computer animation can take many forms, but all teachers tend to struggle with a common problem: how do you teach the technical aspects of the discipline and still have time left for the aesthetics of experimentation and creativity? Paradoxically, our field is perceived as innovative and experimental, yet it is already associated with many clichés and conventions. A primary barrier is the computer animation process itself. In computer animation, the mouse, CPU, and software are interjected between the idea and the result. The challenge is to make the interface a partner rather than a barrier in the process. But too often, students begin posing questions that are third-person neuter: Can it do glass? Can it make ripples? Can it do particle systems? When "it" is referred to as the work, students write themselves out of the process completely. This objectification manifests itself in other ways. Students with a clear technical understanding tend to push to make the software do everything all at once. They apply "lens flare" to every image they create. They program "glowing" buttons for interactivity. They ray-trace glass spheres in reflective rooms till all hours. And they model familiar clichés to learn the tricks of the trade. At the other end of the spectrum, those students who are not familiar with the technology painstakingly move through each demo tutorial. Paralyzed by fear, they never deviate from the path.

So how do you control the extremes and focus on the primary task of making artwork, using the computer as a tool for communicating content?

A creative catalyst is needed to balance curriculum with just the right amount of instruction and experimentation. That catalyst must provide a smooth blending of the technical and conceptual aspects of animating on the computer. There are three basic elements of this overall philosophy: a holistic balance between the conceptual and the technical, a strong emphasis on game theory and experimentation, and a concentration on problem-solving strategies. This workshop introduces participants to motivational exercises suitable for time-based artists, discusses course objectives for developing animation concepts, and concludes with a discussion of the critique process for time-based media.

Motivational Exercises

In theater, the concept of representation connotes realism. Stanislavski's method calls upon actors to represent reality based upon motivations culled from personal experience: a logical, rational building of emotion. In 3D computer animation, a distinct paradox emerges regarding the use of the term "representation." The trend toward photorealism in architecture, design, and scientific visualization can overshadow the "presentational" aspects of computer animation as fine art. This trend is as inhibiting as the realist/naturalist movement was to theater at the turn of the century. Even Stanislavski, the man most associated with the realist movement in theater, ultimately rejected the meticulousness of external realism, and urged his actors to strive for an "inner realism."

Participants in the workshop explore the difference between inner realism and representation through body movement and motivational exercises suitable for time-based artists. These exercises are designed to introduce students to basic animation concepts (exaggeration, gesture, timing, and weight) while at the same time demonstrating the importance of "on-your-feet" brainstorming and improvisation. Examples include the "machine." Participants must build a human machine in which each person generates a sound and a body movement that goes with what the other people are doing. Another example is the object/verb analogy, where participants must "act out" an inanimate object executing a random action.

Karen Sullivan
Claudia Cumbie-Jones

Ringling School of Art and Design
2700 North Tamiami Trail
Sarasota, Florida 34234 USA

+1.941.351.5100
ksulliva@rsad.edu
ccjones@rsad.edu

Concept Development Course Objectives

At Ringling School of Art and Design, our concept development class serves as a useful catalyst for the generation of new ideas. Such a class can also reinforce the need for working through animation as a process, rather than expecting it to spring forth from the head into the computer as an already complete entity. In concept development, we have designed assignments to help students explore multiple aspects of computer animation. In addition to the "acting" aspects of computer animation, directing and design aspects need to be addressed, and a history has to connect with a variety of ideas in order to make sense. The history of electronic art connects to 20th century art through numerous movements, particularly surrealism and conceptual art. It has influenced and is influenced by the history of other time-based arts such as performance, video, and film. It is also shaped by a variety of industrial and scientific applications that have influenced both the visual nature of computer animation and the means by which we create it. These make up the visual references that shape animation ideas. Guiding students toward this kind of visual literacy and making them comfortable with the technology in the process is no small task. Here are some sample objectives for concept development assignments:

Understanding metaphor, symbol, semiotics, and icon:

- Develop concept through dreams, stream of consciousness, or games.
- Understand the use of image as metaphor, symbol, and icon.
- Begin to identify what constitutes a good idea and develop flexibility to change ideas.
- Learn brainstorming techniques such as "what if..." making check lists, forced connections, object analogy, abstraction, finding the essence, "AH HA," metamorphosis, hunches, and use what you have.
- Do storyboards and drawings away from the computer to emphasize the importance of generating ideas on paper and to avoid "Photoshop filter cliches."

Manny Trujillo: Metaphor, Symbol, Semiotics, Icon Assignment

Concept Development for Computer Animators Workshop

Understanding synesthesia, the perfect balance between sound and image:

- Students develop concepts using sound as the source for content.
- Students are introduced to the concepts of juxtaposition, environment, gesture, atmosphere, expression, and impression, as opposed to concepts that have a literal interpretation.
- Limit students to using their mouths or bodies to create sound (this introduces them to the practice of “roughing a soundtrack”).
- Students create a balance between sound and image where neither can exist without the other.

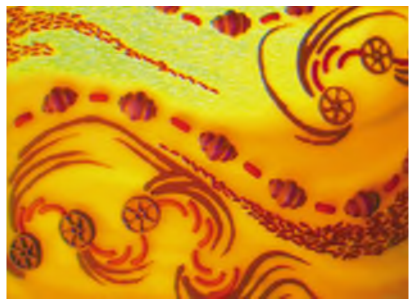

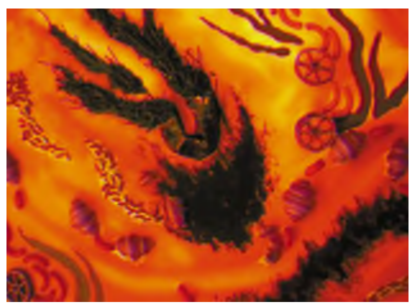

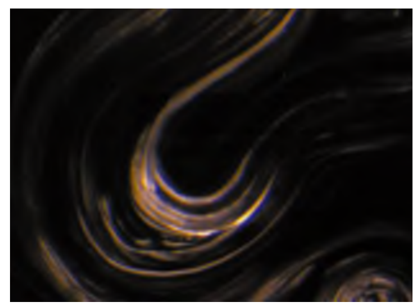

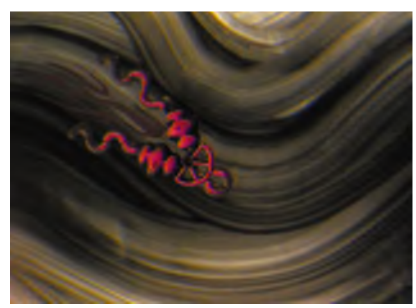

Scott Moore: Synesthesia Assignment

Understanding Transition and framing as language in relationship to layering and compositing:

- Students explore framing as language, developing sutures, sequences, scenarios, and narratives.
- Students explore transition as montage, code, layering, and compositing associated with digital animation.




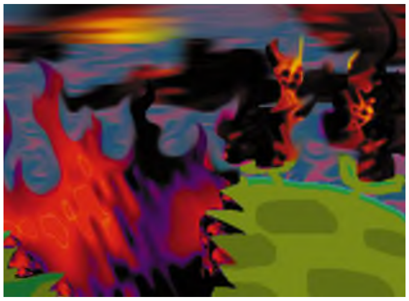

Steve Zafros: Transitions Assignment

Understanding references through research, recognition and juxtaposition:

- Develop concepts through research and comparative analysis. The areas of comparison are visual: between fine arts, time arts, applied arts, and media.
- Include in research an investigation of image as personal, historical, technical, ethical, cultural, to develop a critical perspective.
- Introduce the concept of making a visual comparison over time, as opposed to a verbal or textual comparison, as in a research paper (an extension of the concept of metamorphosis and transition).
- Learn to recognize and avoid cliché, and understand references both in content and technique.

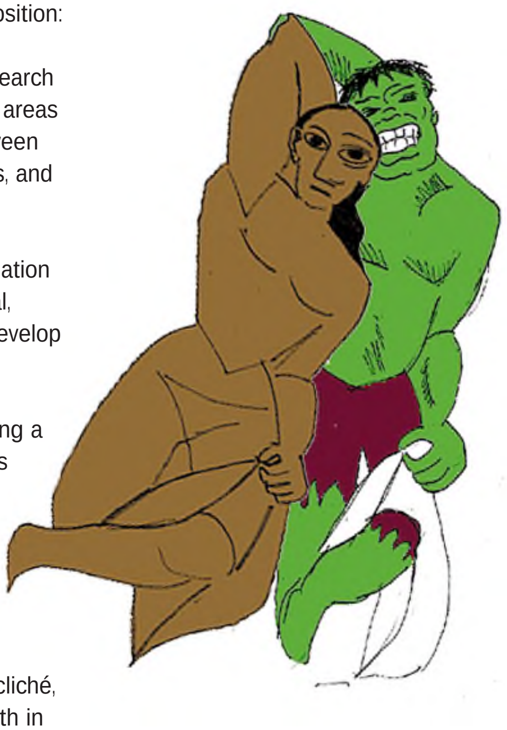

Dani Rosen: Reference, Research and Juxtaposition Assignment

Constructive Criticism for Time-Based Art

In theater, critique is an integral part of the rehearsal process, with a democratic approach to directing. Actors are seldom told what to do by directors; they are coached through a process of exploration and experimentation. There are no wrong ways, only different approaches and interpretations to find the best way.

The animation critique process is similar to the theater critique process. Both are time-based art forms, both require problem-solving strategies where there can be multiple solutions, and both require a "rehearsal" period to test these strategies. For this reason, critique is part of the creative process of animating. In addition, because animation takes so long to produce, a project can easily become stale, and one way to prevent this is to show it to as many people as possible in order to gain a fresh perspective.

Chickscope: An Innovative World Wide Laboratory for K-12 Classrooms

Barbara Mason Fossum
b-fossum@uiuc.edu

"This project involved my students in an area of technology that few high school students have a chance to experience. They learned the possibilities of interactive technology while learning about MRI imaging and fetal development."

High School Teacher

"My students gained knowledge about embryonic development and MRI. They learned new skills in using the World Wide Web and e-mail. My students also felt as though they were a community of learners playing an integral role in a project. They felt like respected people who were given control of an expensive machine."

Middle School Teacher

"My children are constantly involved in discovery learning, always questioning where things come from, what makes things work, etc. Being able to look inside an egg was a wonderful way to learn about life cycles and what goes on inside."

Primary School Teacher

"The children were thrilled to watch the images appear as they manipulated the MRI from our classroom. We watched in wonder as we viewed the pictures of the chick's development inside the egg, and candled our own eggs to see if we could find those same characteristics developing."

Primary School Teacher

Spring 1996 Chickscope: The Beginning

In the spring of 1996, 10 classrooms ranging from kindergarten through high school participated in Chickscope, a collaborative 21-day chick embryology project initiated by the University of Illinois at Urbana-Champaign.

Each classroom was given fertilized eggs, incubators, and educational materials on egg science and candling, courtesy of the 4-H Cooperative Extension Program. At the same time each day, a fertilized egg was placed in a magnetic resonance imaging instrument at the Beckman Institute Magnetic Resonance Imaging Laboratory. Through an interactive Web site, students remotely controlled the MRI device to obtain images of the developing chick. The Web site was developed by the Beckman Institute Visualization Facility, the Magnetic Resonance Imaging Laboratory, and the National Center for Supercomputing Applications.

The classrooms acquired daily images of the egg and shared their observations, predictions, and questions via the WWW. Through the Chickscope Web site, they received educational material, daily reports of the development of the chick, and sample MRI images for the day.

The goal for the Chickscope project was not only to provide students and teachers with access to the MRI system over the Internet, but also to provide them with the supporting infrastructure that is usually reserved for scientists. We realized our goal when the students and teachers became part of the scientific community by sharing their learning exercises, observations, predictions, and questions.

Access to remote scientific instruments using basic Internet tools from the classrooms offers opportunities for K-12 students and teachers to participate in collaborative research and data analysis. We wanted to explore how this project could be integrated in K-12 classrooms in light of the current science education reform initiatives that recommend the use of the Internet for learning and teaching (e.g., Hunter, 1995; Linn, 1992).

Based on student and teacher feedback, the project was well received, particularly in lower grades, because this was the students' first introduction to doing science on the Internet for an extended period (e.g., at least 21 days). The most successful experiences occurred when teachers creatively included Chickscope in their daily lessons. Access to unique scientific resources and expertise provided the students with motivation for learning science and stimulated interest in the scientific enterprise. A surprising result of this project is the continuing use of the Web materials by the participating classrooms (as well as by classrooms that did not originally participate in the project, or had access to the MRI system remotely).

Lessons learned from the Chickscope project relating to the Internet include:

- Students were more involved in Chickscope when it was well integrated into the classroom curriculum plans. The teachers played a critical role in integrating both Web-based and non-Web-based activities into their curriculum.
- Students working in groups were able to share computers and limited MRI time effectively to do serious Internet-based science for an extended period.
- In spite of the complexity of the technology, students and teachers across K-12 grade levels and settings (public, private, home school, after-school science club) benefited.
- After the initial interactive Chickscope project came to a conclusion, the site remained a useful resource in the classroom, using the collected database of images from the original project. The remote instrumentation portion was not essential to the continued use of the site.

C. Bertram Bruce
chip@uiuc.edu

Umesh Thakkar
uthakkar@ncsa.uiuc.edu

Clinton S. Potter
Daniel E. Weber
Janet Sinn-Hanlon
University of Illinois at Urbana-Champaign

Fall 1997
The Chickscope 1.5 Development Team: Interdisciplinary Chicken Soup

The next step required further development of curriculum materials and content that would be well integrated into the classroom curriculum plans. To that end, we recruited the help of the teachers involved in the first phase of the Chickscope project to act as co-developers for the new Chickscope site. Students and teachers learned much about how to collect and analyze data, how to ask questions, and how to communicate their findings with others. They also had opportunities to interact with experts from several disciplines, such as MR imaging, developmental biology, curriculum and instruction, and computer science.

The new Chickscope site includes content and lessons for integrating image processing, biology, chemistry, and mathematics-related materials into the curriculum. Faculty from the College of Veterinary Medicine developed chick embryology units. Faculty from the College of Education developed staff development units and reviewed the design and implementation. The mathematics units were developed by faculty and students in the Mathematics Department at the University of Illinois, using the theme of the egg as a loose framework for presenting mathematical concepts of all types. Researchers from Arizona State University's Image Processing for Teachers program are developing units for image processing and image analysis based on the MRI data sets collected in the first Chickscope implementation. Researchers at the Beckman Institute developed chemistry and biology materials.

Spring 1998
Illinois Chickscope:
Professional Development for Teachers
www.ed.uiuc.edu/facstaff/chip/Projects/Chickscope/ccc.html

Early analyses tell us that Chickscope was successful in terms of immersing students and teachers in a small scientific community (Bruce et al., 1997). Although remote scientific instrumentation is today an exotic and expensive technology for schools, it is already part of the daily practice in research and industry (e.g., Kouzes et al., 1996). This suggests that K-12 students and teachers may need to learn more about this new technology for doing science, and that it is likely to become more commonplace and less costly in the future. This has happened with electronic mail, which is now part of everyday activity in many schools. The particular instruments and scientific domains may change, but understanding the principles underlying this mode of learning through projects like Chickscope should be generalizable to other domains involving new technologies.

The Illinois State Chickscope Proposal was funded by the Illinois State Board of Higher Education to support the participation of 30 Champaign County K-12 teachers in a professional development training program during the spring, summer, and fall semesters in 1998. The project hopes to demonstrate a capacity for sustainable systemic improvement in mathematics and science education. Participating teachers will actively collaborate with 120 pre-service teachers from the College of Education, and with faculty and staff members from several disciplines.
The Illinois State Chickscope project meets both recommendations for professional development as stated in the State Goals for Illinois Learning Technology (ISBE, 1997, p. 19):

- Develop knowledgeable educators to "establish a student-centered, technology-enriched learning environment."
- Require technology-enriched teacher preparation to "ensure that new teachers are prepared to take full advantage of the learning potential of technology."

Previous Chickscope evaluation has shown that it had a positive impact on teachers' ability to promote mathematics, science, and technology learning. As we expand to a county-wide implementation, it is crucial to evaluate what works and why in this proposed project. Specific evaluation questions will focus on the six objectives: What is required to scale up? Is collaboration among teachers, pre-service teachers, and scientists supported? How well does the information infrastructure work? Is there a useful clearinghouse for standards-based materials? Are teachers supported in inquiry-based teaching? Is the project sustainable?

Spring 1998
Chickscope 1.5 Revealed

The Chickscope 1.5 Web site was designed to be an interactive web book, providing content and activities for K-12 students in various curricula. Along the left margin of every page are links that facilitate access to the major sections of the site. These sections ("Explore," "Roosts," "Schools," and "Help") are described in detail below. Along the bottom margin of every page are links to subsections or subunits that deal specifically with the major section visited.

The home page provides a brief overview of the project, the "Yolk of the Day" (a collection of the worst chicken jokes we could find), "Hot off the Press" (any new additions or changes), and "Serious Fun" (a brief description of an activity or interesting page). The page uses server-side technology to create a dynamic and engaging first impression of the Chickscope project. The "Yolk of the Day" and "Serious Fun" sections invoke two PERL scripts that randomly select and return jokes and activities, thereby changing the appearance of the home page upon each visit.

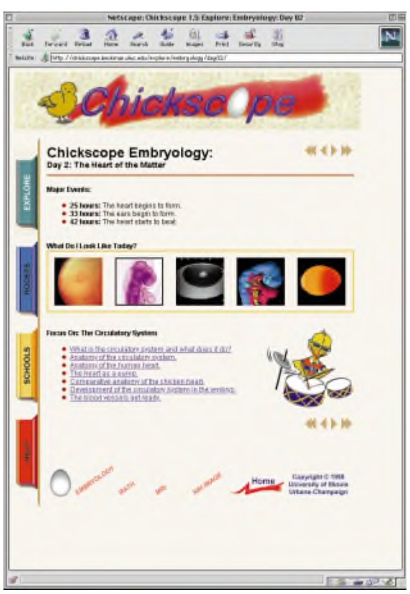

Explore: A Menu of Many Courses

Inside "Explore" one can read about several subunits. One subunit describes the Embryology unit, where students can access the development of the chick embryo day by day. Each day allows students to examine that stage of development in several modalities. Image icons representing these modalities allow the student to look at MRI images, colorful illustrations, images of candled eggs, histology sections, or 3D reconstructed volumetric renderings. Each mode is annotated and contains text related to the developmental stage or an explanation of how the image was acquired. Each day also focuses on a specific element of development and includes lessons or activities for the students. Other subunits include Math, an NIH image tutorial, a tutorial on biological imaging, and an image processing unit developed by the Image Processing for Teachers group at the University of Arizona. The Egg Math pages include Java applets explaining math concepts and exercises to do in class. Topics such as the "ham sandwich theorem" and "egg calculus" engage students in leaning math concepts. Along the bottom of the page are links to these units.

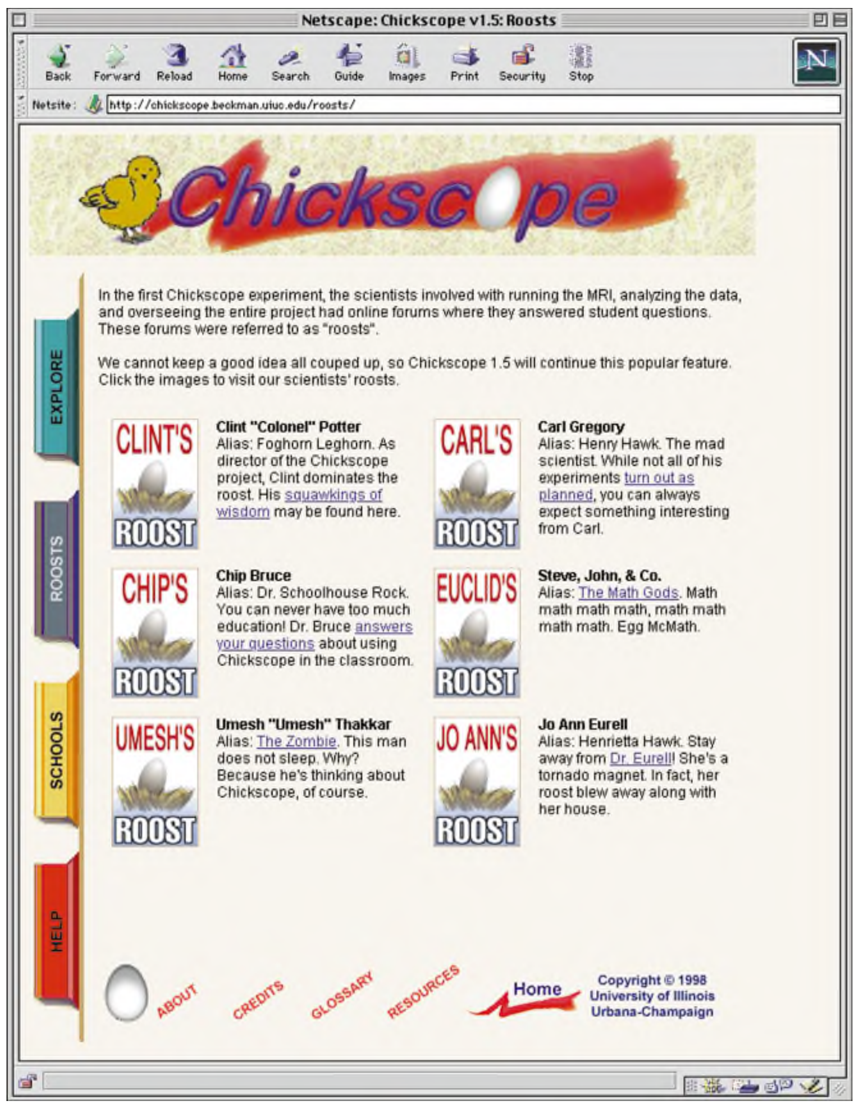

Roosts: Where Inquiring Minds Ought to Go

The roosts provide a venue for the students to communicate with the researchers. Students can email individual researchers at their roosts and ask questions or exchange ideas. Each senior researcher associated with the project supported a Web page for their area of expertise, such as Carl's Roost, which was developed by Carl Gregory, an MRI specialist. Frequently, these roosts were loosely constructed pages put together by the scientists involved with the project. For Chickscope 1.5, many of the roosts became the basis for activities and lessons found in the Explore section. The roosts are used to document project results and respond to students and teachers. Annotated images with notes on the daily growth of the chick are also recorded here.

Schools: The Roots of the Project

The Schools section provides a place for each school to publish its own Web page to report progress in the project. Each school may publish a page providing information on individual experiments they may be attempting as well as contact information for other schools to communicate with them about their projects. This is the place for each participant to establish a unique ownership or presence in the project. Through links on this page, the teachers and the students can communicate with each other via HyperNews, a product from the National Center for Supercomputing Applications that allows teachers and students to post comments using a Web browser and gives greater flexibility in communicating.

Help: "I need somebody..."

The Help section provides online guidance with navigating the site and finding resources. The site includes downloadable software for several of the activities, and the help section covers some of the technical questions that might arise. These downloadable packages are located under the Resources section of the site. Help also provides information on how to use the database, how to use HyperNews, and where to find useful tools.

The Future of Chickscope

What started out as a proof of concept for a remote instrumentation project has become a critical survey of curriculum integration and information design for the K-12 classroom. Chickscope 1.5 opens avenues for many design elements and tools to be used as a framework for delivering content and curriculum to K-12 classrooms. What remains is to demonstrate that Chickscope can scale to a larger national community and to implement this style of delivery to other areas that nurture interdisciplinary study. Chickscope opens the door to bringing new tools and technology to support inquiry-based learning into the classroom via the World Wide Web.

References:

Bruce, B. C., & Thakkar, U. (1997, October). Inquiry-based Learning and Teaching With Emerging Technologies for K-12 Mathematics and Science Teachers. Proposal to the State of Illinois Board of Higher Education.

Bruce, B. C., Carragher, B. O., Damon, B. M., Dawson, M. J., Eurell, J. A., Gregory, C. D., Lauterbur, P. C., Marjanovic, M. M., Mason-Fossum, B., Morris, H. D., Potter, C. S., & Thakkar, U. (1997). Chickscope: An interactive MRI classroom curriculum innovation for K-12. Computers and Education Journal, 29(2), 73-87.

Linn, M. C. (1992). Science education reform: Building on the research base. Journal of Research in Science Teaching, 29, (8), 821-840.

Linn, M. C., diSessa, A., Pea, R. D., & Songer, N. B. (1994). Can research on science learning and instruction inform standards for science education? Journal of Science Education and Technology, 3, (1), 7-15.

Mason-Fossum, B., & Thakkar, U. (1997). Primary school classroom and Chickscope: Studying the egg in the classroom and using the Internet. Proceedings of the 3rd Conference on the Human Factors and the Web, Designing for the Web: Practices and Reflections, Denver, Colorado, June 12.

National Research Council. (1996). National Science Education Standards. Washington, DC: National Academy Press.

Bruce, B. C., & Rubin, A. (1993). Electronic quills: A situated evaluation of using computers for writing in classrooms. Hillsdale, NJ: Erlbaum.

Bruce, B. C., Bruce, S. P., Conrad, R. L., & Huang, H-J. (1997). University science students as curriculum planners, teachers, and role models in elementary school classrooms. Journal of Research in Science Teaching, 34, (1), 69-88.

How to Get Web Presents! Designing a Collaborative K-12 Web Project

June Julian
New York University

Since I first launched the collaborative ecology project, "A World Community of Old Trees," on the Web two years ago, I have received hundreds of contributions from children and adults from all over the world. Each time I opened an electronic file, I felt like I was opening a very special present! What wonderful old tree would be inside?

Using "A World Community of Old Trees" as an example, this paper presents methods for K-12 teachers and students to design, maintain, and evaluate their own collaborative project for the WWW. For me, having a global Web project has been like hosting a party and inviting the world to join in! The purpose of this presentation is to share the excitement of the definite potential of the Web medium for global collaborative learning.

"A World Community of Old Trees" was specifically designed to provide an open digital space for the global community to identify, write about, and document with visual images, the most extraordinary trees in their environment. The project contains three major components: the Tree Gallery, with both scanned art and Web-specific imagery and accompanying descriptive text, the Tree Museum, where references to extraordinary trees are listed in the continually growing Print Sources and Web Sources sections, and Tree Talk, which contains tree ecology facts, personal narratives, photos, participants' responses, and a built-in Comment and Survey Form about the project.

For me the most wonderful section of the Tree Gallery is the Student Projects section. It showcases the wonderful art work, poetry, and stories from children, all lovingly documenting the most extraordinary trees in their worlds. It was especially exciting when these files came into my email account! When I opened them, I was greeted with beautiful tree art works, stories, and photos of children proudly standing next to their trees. Children from many parts of the world, including the United States, Canada, Lithuania, Australia, and Japan, have sent in their tree gifts for the world to share. When visitors responded to a particular student's contribution, I pasted that email directly on the bottom of their page. For example, Rosemary, an eighth grader in Australia, did research on the oldest species of tree in the world, the Wollemi Pine, and contributed both text and photos to the project. When I received an email comment on Rosemary's work, it was added to her page.

Another exciting student contribution to the project was the Interactive Tree Chaos Series page by Josh, a twelfth grader in New Jersey. In this project within a project, Josh invited visitors to interact with his imagery with either conventional art media or with computer graphics and to send their pieces back to the site. His pages include the interesting interactions that came in from several diverse participants: a sound file by an eight-year-old Toronto girl singing "I Am a Tree," a series of computer-graphics manipulations of his imagery by a graduate student in New York City, and similar images from a handicapped artist in the Pacific Northwest.

This experience convinces me that with an open, interactive design, any educational theme that has its basis in the concerns and issues of young people would be successful in an international, shared resource. From my work, I have identified the following factors as considerations for effective design of collaborative K-12 Web projects: interaction, participation, content, promotion, information collection and distribution, and evaluation.

"A World Community of Old Trees" demonstrates several modes of interaction with the project: participants can send in comments on a Comment Form, complete a Survey Form, send in email on tree topics, or interact with existing work already on the site. In one case, they are even invited to manipulate a student artist's imagery. I believe that in creative circumstances, the possibilities for interaction are numerous and fantastic. By its nature, the World Wide Web invites interactions. Thoughtfully designed collaborative Web projects can maximize this innate characteristic of the medium and celebrate it within any thematic context.

Within its three major components, TREE GALLERY, TREE MUSEUM, and TREE TALK, global participants of all ages rejoice in the planet's most extraordinary trees. I am able to witness the pictures and words from many hearts. By making participation in a collaborative project as open as possible, project designers can maximize the potential of the Web for engaging a large and diverse group.

With regard to the content of "A World Community of Old Trees," my purpose was not to examine global tree ecology problems nor to suggest solutions to them. Rather, my aim was to provide a framework for participants to have relationships with trees, and through an interactive digital context, also with each other. I was more interested in seeing what came in rather than in presenting any particular view. Initially, I presented a few demonstration pages featuring my own pieces in a variety of media as examples of page layout and possible style. As each contribution came in to the project, I put it on the site right away for potential participants to see, and to add to the ongoing demonstration of possibilities of interpretation. Content design should be inclusive and open, within a flexible structure.

Since "A World Community of Old Trees" is a Web project, the natural dynamics of the medium itself function quite spontaneously for the project's promotion. Since it is developing over time while open to public view, it is its own best advertisement. Still, initially, I had to consistently publicize the project by repeated email postings on numerous subject-matter-related listservs. After a while, though, the project picked up its own momentum, and began to have a life of its own, gathering participants daily. Information will reside silently on a server unless determined efforts are made to promote those links to the target audience. A very effective method to consider is the use of complementary links.

With regard to information collection and distribution, I found that the practice of building the site as the contributions came in was very effective, particularly for students who were anxious to see their work on the Web. As soon as I received them, I put new files on the server, and immediately sent out email messages that included the new URL.

Within a Web-based collaborative project a lot of teaching goes on in the most unexpected and marvelous ways, but how can you best evaluate the project? Once again, the interactive nature of the Web itself can provide the means. As with "A World Community of Old Trees," various fill-out forms can be placed on the site requesting input on the the project's objectives. And, if your email inbox is always full of new messages from people all over the world, discussing, suggesting, questioning, and telling, you will be well-supplied with feedback on your project.

But most of all, it is the quantity and quality of those precious gift files that speak to the strength of your Web project. How do you get the best Web presents? If you carefully consider the factors of interaction, participation, content, promotion, information collection and distribution, and evaluation when you begin to design your collaborative Web project, you have the best chance, I've found, of maximizing the potential of the Web, an entity as vast and synergistic as the human spirit itself.

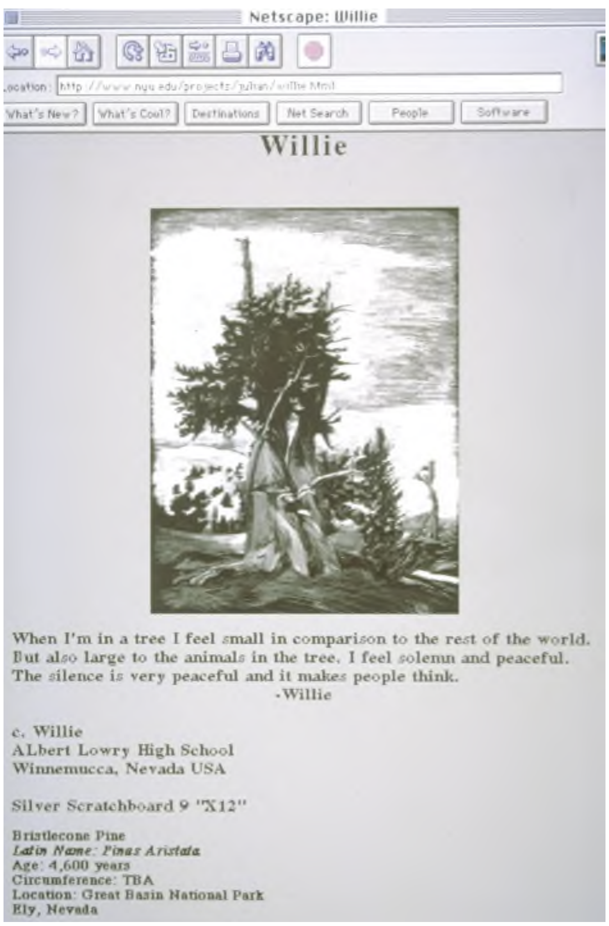

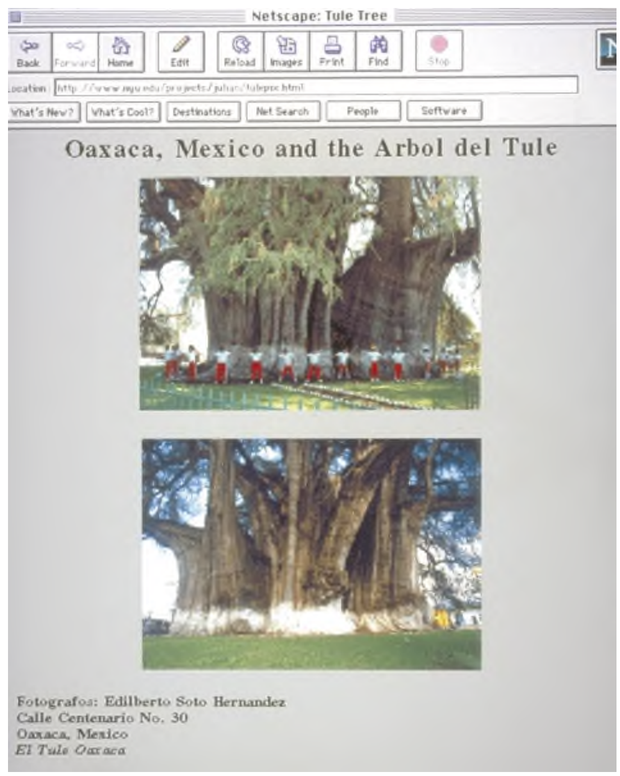

Improving Instruction and Staff Development by Building K-12/University Partnerships

Rodney Zagury
Adjunct Instructor
New York Institute of Technology
bbz@erols.com

Artists have long been the voice of society and civilization. Whether creating art for the masses, for the select few or for themselves, they have kept a finger on the pulse of evolution. Computer animators bring life to the imagination with digital manipulation and personal observation from high up on the wave of the Zeitgeist, by producing art that has strong didactic and entertainment potential.

As a computer animator working in three-dimensional graphics, one of us draws on personal experience to give clients an interpretation of the concept they have developed for a particular project, client, or design. As a teacher of computer animation, he gives his students a peek into a finite world that allows for an incalculable number of personal interpretations of any single idea or vision. Most of those students are seeking a Master of Arts degree in communication arts, specializing in computer graphics for advertising design and animation, or in television production. Many of them have little or no experience with computers; most are traditional artists. With those students, a good place to start is by having them tap into their own experiences and resources to develop a personalized approach to 3D graphics, without letting the complexity and technical demands of this medium discourage them from pursuing a career in animation. The most challenging aspect of the job is the varied levels of computer literacy and 3D theory encountered in those classes.

What has made the job of college instructor easier is that the growing use of computer animation for education, science, business, medicine, law, and entertainment has widened the playing field for the animators of the future. People from all walks of life are likely to find something that will interest and occupy them in computer animation. Far from dehumanizing, the computer has the power to bring the individual closer to the life rhythms that fuel the artist's vision.

Computer graphics has truly become an artist's medium; it is now a matter of making these new tools available to those who would benefit most from using them. As an introduction for these university-level students, we like to tell them that they have all been prepared to work in 3D computer animation, because they live in a three-dimensional world. They are asked to think of tasks they might have performed recently, or observations they might recollect with any sort of detail. Have they looked, but really looked, at the way light reflects off an apple, or the way sunlight enters through a window? One student actually said that she had no practical experience that relates to 3D modeling, but she had been working as a recreational therapist in arts and crafts.

What is most fascinating about working in 3D is that instructors are called upon to connect with and use almost everything they have learned and experienced in an entire lifetime. Some connections are obvious and quite practical: photography, ceramics, sculpture, fabric design, dance, and music. Some are less obvious: watching television, reading, looking up at clouds, cutting up a pear (while observing it carefully) and just living and soaking up the environment. It is good to think that everything is useful because 3D is so all encompassing. A recent project had one of us digging into memories of observing abdominal surgery because the client wanted realistic "living" organs: a heart, a stomach, and intestines. Photographs were a big help, but what helped most were recollections of light, translucent and reflective tissue, and true, living color.

Today, 3D software packages are available for the PC, and they range in price from the very affordable to the expensive, but not quite as expensive as they were five or 10 years ago. That affordability was the impetus for this paper. We wondered about ways in which university and K-12 educators might collaborate to take advantage of all that is now available. Many elementary and high school students are actually already experimenting with computer graphics and 3D art; the animators of the future are already training and preparing themselves for this growing industry. As educators at these different levels, we can all work together to guide and motivate students who have an interest in the field and want to acquire practical skills that could further their goals as future animators. Besides the obvious – skills in traditional animation, computer science, film and video – there are a number of extremely useful endeavors that would assist students in their goals of entering this field. These are just a few of them: dramatic arts (acting, directing, lighting, set design), performing arts (dance, music, opera), writing, storytelling, studies in folklore and mythology, photography, yearbook layout and illustration, drawing, painting, color

Brian Monahan
Director of Information Systems
North Rockland Central Schools
65 Chapel Street
Garnerville, New York 10923 USA
bmonahan@ric.lhric.org

theory, sculpture, ceramics science, physical and biological sciences, history, art history, architecture, mathematics/geometry, computer-related and business-related communication, sports and physical activity in general, television and film production.

A student can derive a great deal of pleasure from a project that focuses on personal interests. A teacher can get a lot of mileage out of that student with a well-directed project, and a lot can be learned by both of them in collaboration. One of us has strong memories of a dialogue written for high school Spanish that would easily translate into a good animated short. It is based on a fable by La Fontaine.

How does this partnership work? The North Rockland Central School District and New York Institute of Technology are near one another but are certainly not next door to one another. NYIT has several campuses, but the closest is at least a 90-minute commute from the school district. So, the partnership was established using the facilities of the district. Fortunately, North Rockland is a public school district with a strong commitment to technology and the arts. The district, which has 7,500 students, commits over $1.5 million per year to technology, creating state-of-the art computer laboratories, placing computers in individual classrooms, and putting time and money into staff development. The commitment to staff development is also reflected in the presence of a busy Teachers' Center. It is the Teachers' Center that manages the program that NYIT offers in Rockland, taking care of the mundane details of registration, room assignments, etc. The Teachers' Center is located in North Rockland, but it also provides staff development for eight other school districts, thereby increasing the pool of students.

A consideration of such a partnership is that it must operate within the context of the needs of the participating teachers. There are many teachers in the region who are anxious to explore the use of computer graphics. However, the reality of the certification rules in New York and New Jersey (the areas served by the partnership) is that most teachers cannot achieve certification with degrees in computer graphics. What we have done through our partnership is incorporate significant components of computer graphics into education courses that meet the state standards for certification. Teachers in our courses explore successful models of the use of computer graphics in schools. As an example, as part of a computer graphics component of our Introduction to Educational Technology course, one of our students explored art education in traditional classes (Dana 1993) and in gifted and talented classes (Banbury 1996). One of the ways we have made those courses successful is through an annual Summer Institute in Educational Computing, which the North Rockland School district runs with the support and assistance of NYIT. That institute gives K-12 faculty an opportunity to explore a number of areas relating technology to education. Computer graphics is always one of the more popular.

What impact has the partnership had on our students? Because of the training that our faculty have received in a variety of courses in the NYIT partnership, our teachers are using computer graphics at many levels with their students. At the elementary and middle school levels, teachers uses programs as simple as Microsoft Paint and Print Shop to teach basic concepts in art and design. At the secondary level, teachers across many curriculum areas use software such as Photoshop, AutoCad, and Corel Draw to help students learn concepts of art, design, and even engineering.

As we contemplate the many digital microcosms of 3D graphics, we think of Vincente Huidobro, a Dada and "Creationist" poet. His aesthetic theory for Creationism proposed to "create a poem the way Nature makes a tree _ make real that which does not exist ... say things which [without you] will never be said," and ultimately stated that "the poet is a small God."

References

Dana, Ann. Introduction of Technology into the Art Curriculum. ERIC Document Reproduction Service, ED370594, 1993.

Banbury, Mary. Computer Art Programming for Gifted and Talented Students. Gifted Child Today, January 1996.

The Physics 2000 Project: Interactive Physics on the World Wide Web

Martin V. Goldman
David Rea
University of Colorado at Boulder

What is the Physics 2000 Project?

Created by leading researchers, the Physics 2000 Project has generated an innovative interactive Web site[1] designed to make modern physics and technology accessible to K-12 students, college students, and the general public. A key element in the success of this ambitious project has been the use of Java applet technology to create over 35 "virtual experiments" that are controlled using the computer mouse. By combining graphics primitives, concurrent threads, and GUI input events, Java permits rich inter-active animations to run locally on Web clients.

Users of Java-enabled Web browsers such as Netscape 3.0 find their screens transformed into a laboratory containing experiments ranging from an X-ray fluoroscope[2] and CAT-scans[3] to interactive laser and magnetic-evaporative cooling[4] for creation of the Bose-Einstein condensate. This "virtual" laboratory has certain advantages over real laboratories. It allows the user to see the physics concepts as they are visualized by leading scientists. These dynamical visualizations convey far more than static text, or even a real experiment. Interactivity further enhances the learning experience.

For example, when charges are shaken (using the mouse), the users see how the electric lines of force wiggle[5] and an electromagnetic wave is launched. The frequency can be varied and the user notes that this affects the wavelength. An option to derive the algebraic relation among wave speed, wavelength, and frequency is offered. In another applet, users shine different colors of light on an atom and observe how the electrons change in energy and position. These and other interactive virtual experiments are demonstrated in this paper.

The interactive demonstrations are unusually tactile and responsive, as well as just plain fun. This makes the science as concrete and intuitive to the student as it is to the seasoned researchers who designed the site. The subject matter is primarily 20th century science and technology. Devices such as X-ray machines and microwave ovens lead into explanations of electromagnetic waves and how they interact with matter. Recent exciting research discoveries such as Bose-Einstein condensation and novel quantum interference are also featured.

In addition to virtual experiments, other pedagogical strategies are used to enhance the presentation. Dialogue between cartoon characters provides exposition into which the virtual experiments are embedded. By proceeding from familiar technological devices to the more abstract underlying concepts, the physics is made relevant to the non-scientist. Various levels of complexity are incorporated through the hypertext format so that users can easily select the level that matches their background and interests. The cartoon characters serve to identify the level of difficulty of the subject matter. Workshops with high school teachers have been employed to evaluate these strategies and useful feedback was incorporated into the Web site.

The Web site has received considerable recognition. It has been endorsed by the American Physical Society and the Exploratorium in San Francisco, and it has drawn favorable responses from educators, students, and users from around the world. Future development will include units on lasers, compact disc players, microwave ovens, molecular bonds and basic chemical processes, particle accelerators, thermonuclear fusion, and many more.

The innovative form of presentation transcends the traditional divisions of "elementary" and "advanced" concepts and topics, which form the foundation upon which most physics curricula are built. As it expands to cover more material, this project should be a valuable educational resource for science teachers and students at every level.

How Did the Physics 2000 Project Get Funded and Managed?

The Physics 2000 project was originally conceived as an in-class resource to be used in conjunction with innovative undergraduate physics courses for non-majors. Support for the project was initially sought from the undergraduate curriculum division of the National Science Foundation (NSF), but a proposal for seed money was turned down. The focus of the project was then broadened to include outreach, especially to K-12 educators throughout Colorado. A proposal with this new emphasis was sent to the Colorado Commission on Higher Education (CCHE). In addition to development of a Web site, the proposal provided for a sequence of workshops to be held with local and statewide educators. The proposal was approved, and a generous level of support was furnished by CCHE for 18 months. This grant ended June 30, 1998. At present, we are seeking support to continue for another three years. Once again, we are going to NSF, but to a different division. We are also seeking to secure corporate support through the University of Colorado Foundation.

The grant from CCHE enabled us to hire a cadre of artists, designers, cartoonists, HTML formatters, and programmers to create a unique and highly professional technical staff. In addition, graduate students and undergraduates were hired. A group consisting of these personnel together with faculty from the Physics and Chemistry Departments at the University of Colorado met weekly to discuss strategies and review the progress of the Web site. In between these meetings, a student (graduate or undergraduate) would typically meet with a professor and our programmer, David Rea, to develop a flow of pages in the Web site devoted to a particular topic. Once the general flow of information was established, both faculty and students wrote dialogue for our cartoon characters. Often a particular Java-applet "virtual experiment" drove the dialogue. In general, dialogue was felt to be the best form of exposition for the Web, since it nicely parses the information into a few sentences at a time that can be easily read on the screen.

What is So Unique About the Java Applets on This Web site?

One of the effects we have striven for in our applets is tactility: the illusion of physical reality in the graphics. Based on both observation of students and personal experience, we have found that simulations that look more real are also more engaging. The computer game industry, obviously, has discovered the same thing.

For the most part, tactility in Physics 2000 is accomplished through entering the third dimension. That is, we attempt wherever possible to step beyond flat, monochromatic graphics and instead use shaded 3D objects, whether in the spheres and arrows we use to represent particles and vectors or the buttons and sliders we use for controls. Liberal use of drop shadows (even for some objects, like electrons, that don't really have shadows) further enhances the illusion.

The major drawback with these techniques is in loading time, as they often require extra images to be imported over the net. Loading times can be minimized with the use of JAR formats, but so far we have avoided this approach because of the scarcity of JAR compatible browsers.

Another important component in achieving tactility is to design applets in which the user interacts with the simulation itself, rather than merely adjusting sliders and pressing buttons. Although this approach is not always practical, we try to use it whenever possible because it makes the objects feel even more real. Furthermore, screen controls almost always imply a finite number of possibilities for the applet, while interaction with the components themselves vastly increases the variability of the simulation. Pedagogically, this leverages the same advantage that a Lego set has over, say, a toy truck or plane. Flexibility encourages experimentation and thus extends engagement.

Java was pretty much the only language choice for Physics 2000, given our requirements for multi-platform compatibility and integration with the Web. However, it is worth mentioning a few of Java's strengths and weaknesses. Perhaps Java's greatest asset is its smooth handling of graphics. It is a relatively simple process to import transparent GIF images and animate them, creating rather complex simulations with minimal effort. In addition, Java makes it extremely easy to convert between images and integer arrays, allowing for a variety of algorithmic image manipulation. Although the language itself supports alpha (transparency) channels, current implementations don't support intermediate alpha values, only complete opacity or transparency. However, the ability to generate non-rectangular images is very useful.

Ignoring for now the bugs in various Java implementations, the single greatest limitation in Java is the inability to graphically step outside the bounds of the applet to draw on the surrounding Web page. In other words, when you define a region for the applet on the Web page, the applet is not only stuck within those confines, but also text or images from the page can not intrude inside that region. Changing this design flaw with the addition of just a few methods in the Java language would not only allow much more visually interesting effects, but would also save significant page real estate. Imagine a model of the atom where the electrons fly out in huge orbits across the rest of the page. (For an example of how this would look, see the applet at:

www.colorado.edu/physics/2000/applets/newhome.html)

The Physics 2000 Project: Interactive Physics on the World Wide Web

Comments from users of the site:

"Your site illustrates the power of Java applets to make complex subjects come alive in a way that visitors, young and old, can understand and appreciate. I hope that you and your team will be able to create more experiences of this quality in the future."

Ramon Lopez, Director, Education and Outreach, The American Physical Society, Washington, DC

"I am a high school physics teacher in Highland, Illinois. This site is terrific. I have been looking for some way to use the computer to enhance the learning process, but I have found that most programs and sites seemed to be unfriendly and inaccessible. This site is fantastic! I can't wait to let my students know about the site and have them begin to explore for themselves. Keep up the good work. Thanks for your hard work!"

Cliff Parker

"I'm a second-year student at the University of Michigan, and I'm currently taking a course in 20th Century Concepts of Space, Time and Matter. My professor put links on his own Web page to your Web page, and I've found it to be extremely helpful. I think it's creative and instructive at the same time. I never realized what a good job cartoon characters could do of teaching Quantum Physics! Thanks for your help!"

Peter Handler

1 www.colorado.edu/physics/2000
2 www.colorado.edu/physics/2000/xray/index.html
3 www.colorado.edu/physics/2000/tomography/x _rays.html
4 www.colorado.edu/physics/2000/bec/index.html
5 www.colorado.edu/physics/2000/waves_particles /wavpart4.html

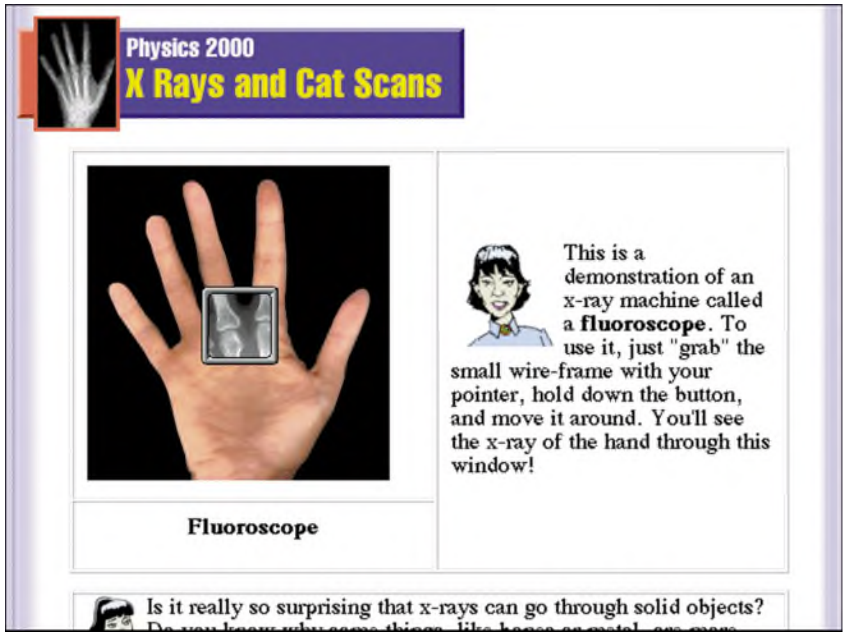

Fluoroscope applet. (A tactile applet for all ages.)

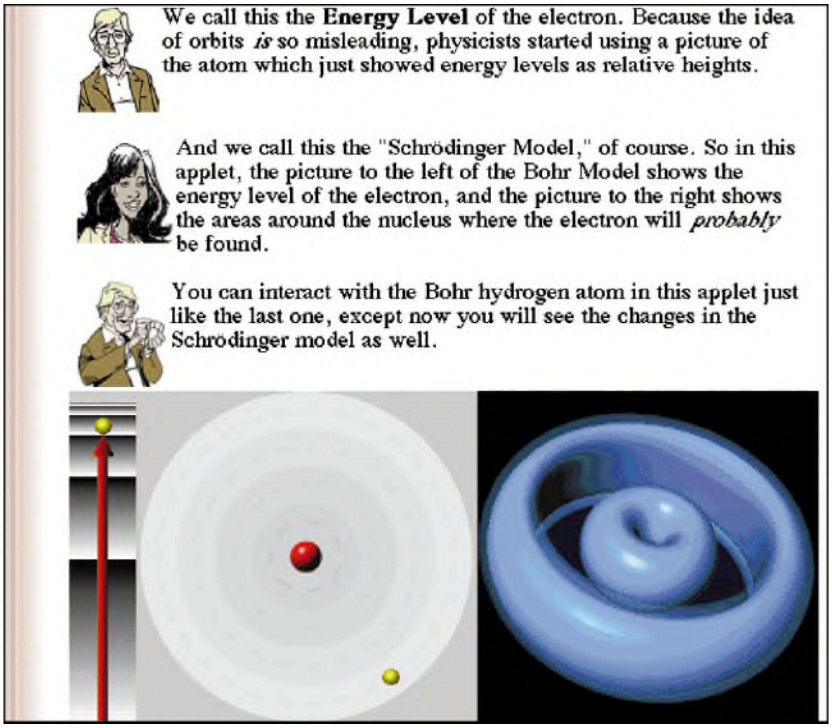

Schroedinger and Bohr atom applet. Click on an orbit to induce a transition involving a photon. (Modern physics applet for advanced high school or undergraduate college physics.)

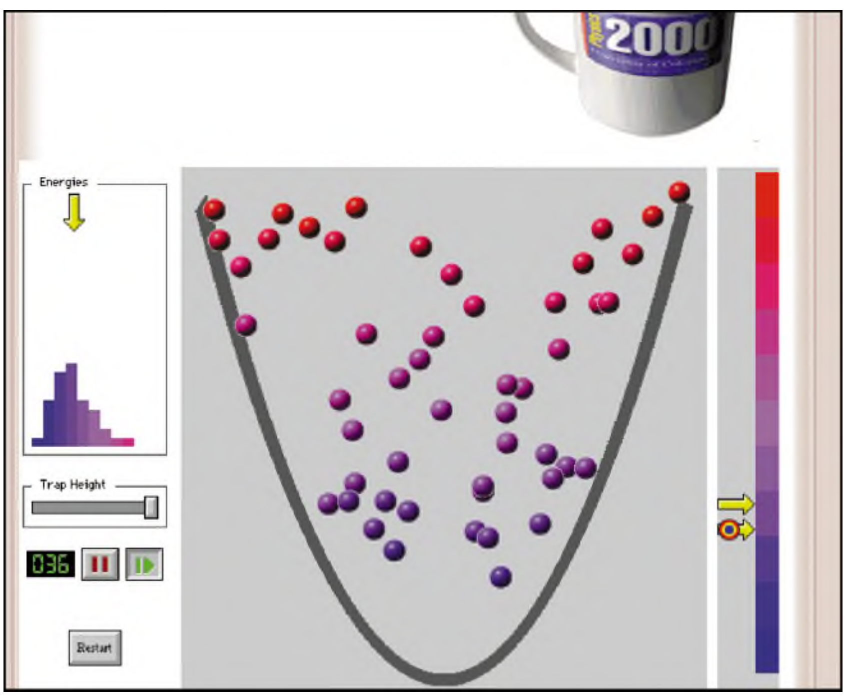

Bose-Einstein condensate applet. Magnetically trapped atoms are evaporatively cooled to low temperatures, producing a new state of matter. Lower the trap height to “evaporate” the fast atoms. (A modern laboratory achievement made understandable to the general public.)

The Ceren Web Resource: Enabling Students to Become Anthropologists In A Virtual Site

Jenniffer Lewin, Mark Ehrhardt, and Mark D. Gross
Sundance Laboratory for Computing in Design and Planning

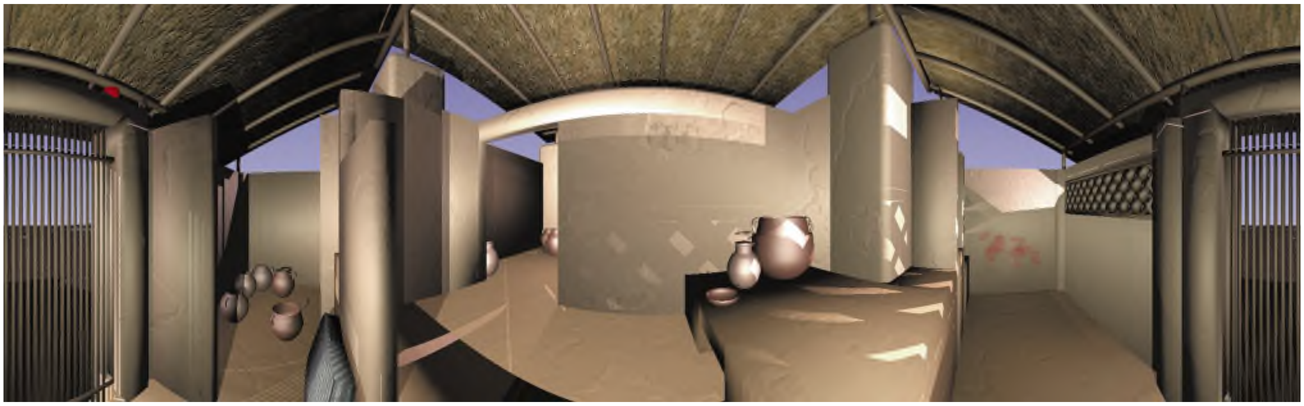

QuickTime VR Panoramic Image from the Ceren Web Site

Using content from the University of Colorado Anthropology Department and our own knowledge of multimedia system development at the College of Architecture and Planning, we are developing interesting, informative, and interactive Web learning resources. Our topic: an anthropology site called Ceren, an ancient agricultural village in western El Salvador buried by volcanic ash over 1,400 years ago. Our goal: to create a virtual anthropology site on the Web, with interactive QuickTime VR, interactive database search tools, image applets, and detailed computer renderings of what the site may have once looked like.

Currently in use by anthropology classes at the University of Colorado, the Ceren Web Resource incorporates an array of Web multimedia technology to link visual images with original excavation text, notes, and discoveries. The Ceren Web resource has attempted to lead the way in exploring the use of hypermedia tools in education. We aim to go beyond a tour of ancient buildings and enable students to begin to think and act like anthropologists. Using the Ceren site, students can participate in the excavation process as on-location anthropologists, putting together the puzzling pieces of what ancient household life in Meso-America was truly like.

The Ceren Village

Discovered in 1978 by anthropologist Payson D. Sheets and registered on the UN Heritage list, Ceren has been acclaimed the “Pompeii of the New World” [Sheets, 1992]. Due to a sudden volcanic eruption, villagers fled the ancient site leaving everything as it had been used in daily life. The volcanic ash prevented decay of all objects. Thus Ceren offers a detailed and exciting view of ancient household life. The Ceren Web site provides a detailed understanding of the ancient structures through computer reconstructions linked to the original excavation text, notes, and data. Used as part of a hypermedia system, computer graphics go beyond static illustrations and become powerful and informative tools for teaching and learning anthropology.

Computer reconstructions of structures 6, 1, and 12 (top to bottom). Each are household living structures at Ceren.

University of Colorado at Denver
Boulder, Colorado USA 80309-0314
ceren.colorado.edu

A Shockwave slide presentation on the Web site showing an excavation image of Structure 12. Structure 12 may have been used to train a female shaman.

Information in Context: Linking Images and Interpretation

Although animation and computer graphics create exciting imagery and help students understand what the ancient structures may have once looked like, without a strong connection between an image and the field data that the image is derived from, students cannot fully understand the site. The Ceren Web site presents images in context.

For example, using QuickTime VR and interactive movies, students can navigate an interpretation of one of Ceren's many structures. Each movie presents numerous artifacts. When selected, each artifact displays text including the artifact type, specific excavation notes, dimensions, and an excavation photograph.

If students decide to study a particular subject or artifact further, they need not leave the image and movie they are in. Pop-up applets allow students to enter a discussion group, post questions, search over 300 pages of online excavation notes and reports, search an image database, or view an interactive slide show on the structure.

Clues Guide Learning Exploration

Each page in the Ceren Web site contains an instructor comment area for clues and questions that prompt students to explore a particular subject. For example, viewing the QuickTime VR or a movie file for household 1, a student finds the following questions from the instructor:

"It was a surprise to us to find that household 1 had more than 70 ceramic vessels. Do you know of any households, in your experience, with that many containers? What is the range of uses to which ceramic vessels can be put?"

Answering this question requires a student first to browse the site, looking at interpretations and arguments made by professional anthropologists who have written about the site. This process, which requires students to do more work than if they merely looked up the answer in their textbook, engages them more seriously in questions of anthropology.

Pilot Study of Learning Effectiveness

Our pilot studies with students of anthropology and architecture revealed several advantages to the Web site over the traditional textbook. First, students are more interested in using the system. Students preferred the system for both its graphical interface (the information was easier to understand when used in conjunction with interactive images), and second for its accessibility. Students could use the site at home and access excavation data that were in most cases difficult to find. Because anthropologists use the Ceren Web site to create archives of their current research, the Web site contains the most comprehensive and up-to-date material on Ceren.

Second, students who used the Ceren Web site demonstrated a greater understanding of the spatial and architectural attributes of Ceren.

Making it Easy for Teachers to Author the Web Site

The Ceren Web site allows teachers to easily change the content of the site by filling in forms online. Thus teachers can update textual information, add new notes, add new artifacts, or add homework assignments. While students view an interactive site with movies and textual links, instructor with almost no technical expertise can log on to a back-end version to add information, update, or make changes. This allows a teacher to change the system daily to fit class needs, and avoids limiting site design to Web developers.

Reference

P.D. Sheets, 1992. The Ceren Site: A Prehistoric Village Buried By Volcanic Ash in Central America. Harcourt Brace College Publishers, Fort Worth.

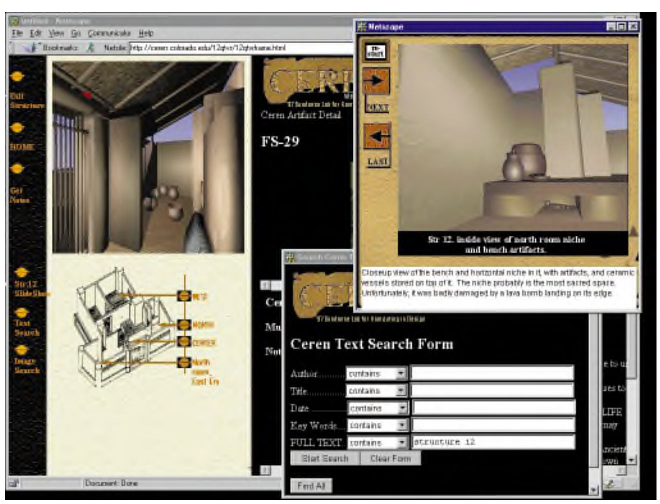

Top left: A QTVR movie showing a different portion of the room

Below: The map of Structure 12

Top right: A pop-up slideshow on Structure 12, with notes from instructor

Below: A pop-up text search window

An Interactive Course on Fractals and Chaos

Cary Laxer
Department of Computer Science

Aaron Klebanoff
Department of Mathematics

This paper describes a course that teaches fractals and chaos via significant student interaction with computer software. The course was developed and taught jointly by the authors (a computer science faculty member with expertise in computer graphics and a mathematics faculty member with expertise in chaotic dynamical systems). Exercises for students to investigate properties of chaotic dynamical systems were developed for computer algebra systems (Maple and Mathematica). Interactive computer-graphics-based software developed by the instructors provided the ability for students to investigate fractals. Lecture material provided the theory on fractals and chaos and the relationship between them. Project work allowed the students to reinforce their understanding of the concepts. Students learned both mathematical and computer graphics concepts in the course, and could receive course credit either in mathematics or computer science.

Introduction

Ever since Mandelbrot popularized fractals, students of mathematics have been intrigued by the theory behind them. Likewise, students of computer science, and especially computer graphics, have been engrossed in the beauty of fractal images. It wasn't long before students requested courses to learn about fractals and chaos, from both mathematical and computer graphics perspectives.

For several years, each of us has been teaching special topics courses on this material within our respective departments. The computer science course covered fractals from a computer graphics perspective but included some mathematics for theory. On the other hand, the mathematics course was mostly theoretical, using some simple computer graphics programs to demonstrate fractals as examples of chaotic dynamical systems. In our discussions, it became apparent that the two courses could benefit from each other. Thus, we decided to develop a single course covering fractals and chaotic dynamical systems that we would teach together. The course would be cross listed in both departments, and students could earn either computer science or mathematics credit for the course.

The Course

We each received a one-course release in the winter term of the 1996-1997 school year to develop the syllabus and class materials for the new course, which was taught for the first time in the spring 1997 quarter. As a unifying theme, we chose to concentrate on dynamics of the quadratic map ($z \to z^2 + c$). This enabled us to learn about graphical analysis of maps, controlling chaos, Cantor sets, fractal dimensions, the Mandelbrot set, and Julia sets. To add breadth to the course, iterated function systems, random fractals, and other maps were studied as well. We chose "Chaos and Fractals: New Frontiers of Science"[1] as the reference book for the course because of its readability and broad coverage of the discipline.

Our primary goal was to get our students motivated to learn about fractals and chaotic dynamical systems. Since most of the topics are motivated by computer graphics, we developed software for the students to help them understand the course material. This included demonstration software (for observing and exploring properties of fractals and chaos) and software for students to modify as part of an assignment. In addition, we developed explorations with computer algebra systems (both Maple and Mathematica) so that our students could understand the underlying computations involved in creating the fractals. Since the course met in a classroom equipped with computer workstations, these computer algebra explorations could be done during class time in conjunction with the corresponding lecture material.

To allow the students to explore the Mandelbrot set and Julia sets and the relationship between them, we developed an interactive software product called MandelbrotExplore (written in C using OpenGL on SGI workstations). The program displays four windows on the screen (see Figure 1). In the upper left window, the Mandelbrot set initially appears. Students can, through use of the mouse, select a region of the Mandelbrot set to zoom in on and explore. By clicking on a point in the Mandelbrot set window, the Julia set for that point is generated and displayed in the upper right window. Students can zoom in on regions of the Julia set as well. Clicking on a point in the Julia set window generates the orbit for that point and displays the orbit in the lower left window on the screen. The lower right window is intended for showing numerical values for orbits inside the Julia set window, although it is not yet functional. (Source code, instructions on how to compile and use the software, and additional example screen shots can be found on the CD-ROM.) In addition to being a powerful research tool, the program provides computer science students with a complex (no pun intended!) example of how to produce fractals.

It is crucial to have a good comprehension of bifurcations to understand the Feigenbaum diagram as well as the Mandelbrot and Julia sets. To build this comprehension, we spent a fair amount of time exploring bifurcations with a computer algebra system. In addition to computing bifurcation values, the computer algebra system enabled students to find formulas for and graph the curves in the Feigenbaum diagrams.

Rose-Hulman Institute of Technology
Terre Haute, Indiana USA
Cary.Laxer@Rose-Hulman.Edu
Aaron.Klebanoff@Rose-Hulman.Edu

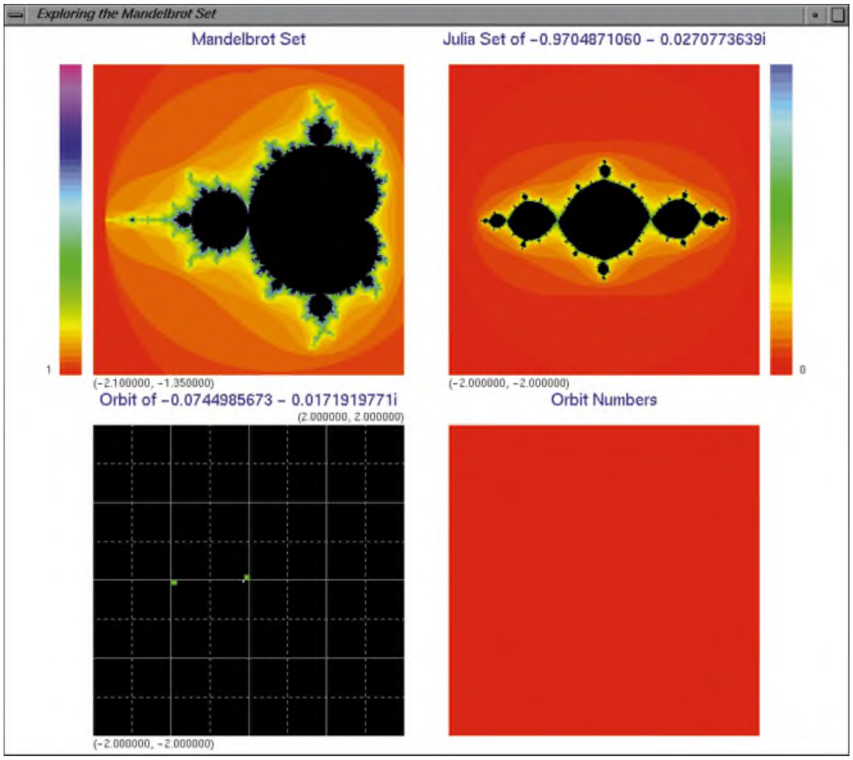

Figure 1 Screen image of the MandelbrotExplore program, an interactive software product that allows the user to explore the Mandelbrot set and Julia sets and the relationship between them.

When studying complex maps, students could repeat numerical computation of complex-valued orbits, and even determine the equation of the largest bulbs of the Mandelbrot set for the quadratic map, as shown in Figure 2. The computer algebra system was also used to understand the action of affine maps to help understand how iterated function systems worked. We also spent time learning how to control the chaotic quadratic map. Figure 3 shows a Maple worksheet in which the unstable period two orbit of the chaotic quadratic map is controlled by the OGY technique.[2]

Results

Enrollment in the first offering of the course was 19 students, consisting of computer science majors, mathematics majors, and a few engineering majors. Students were motivated to learn several major topics in the course, including:

- Complex algebra (through Mandelbrot and Julia Sets)
- Computer graphics
- Iteration (through dynamical systems and feedback)
- Bifurcation theory (through the Feigenbaum diagram and Mandelbrot Set)
- Cantor sets and countability
- Probability (through Brownian motion)

Students especially liked the week spent on controlling chaos. The computer algebra system was especially useful for helping students understand the underlying computations to control chaos. Our time and energy was devoted to the Mandelbrot set for the quadratic maps, but once students realized that there are other Mandelbrot sets (and subsequent families of Julia sets), the graphics-minded students seemed especially interested in working beyond the course to discover the shapes of Mandelbrot sets for other maps.

Conclusion

The course is now being offered annually as a joint elective within both the mathematics and computer science departments. Enrollment for the spring 1998 term is 19 students

References

1 Peitgen, Heinz-Otto; Jürgens, Hartmut; and Saupe, Dietmar. Chaos and Fractals: New Frontiers of Science. Springer-Verlag, 1992.
2 Ott, Edward; Grebogi, Celso; and Yorke, James A. Controlling Chaos. Phys. Rev. Lett. 64:1196-1199, 1990.

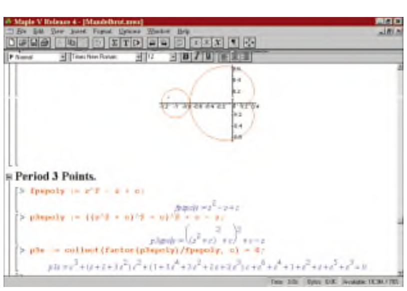

Figure 2
A Maple worksheet showing the equations and graphs for the largest bulbs of the Mandelbrot set.

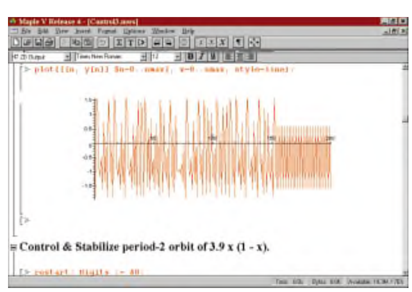

Figure 3
A Maple worksheet showing the unstable period two orbit of the chaotic quadratic map controlled by the OGY technique.

The Art and Science of Multimedia: An Interdisciplinary Approach to Teaching Multimedia at a Liberal Arts College

Abstract

At Wellesley College, very rarely do the Fine Art and Computer Science faculty cross paths. At least that was the case until last year when we taught an experimental course that brought together the work we were doing in our respective corners of multimedia into one class. The course was taught for a second time in the spring of 1998 semester and has been incorporated into the curriculum. This paper describes our experience organizing and teaching such a course.

An Interdisciplinary Course

With the growth of multimedia, the boundaries between traditionally unrelated disciplines have blurred, requiring the collaboration of computer professionals, artists, engineers, and scholars from all fields. Likewise in teaching a multimedia course, we believe that collaboration is an optimal format to thoroughly cover the diverse components of multimedia and to create an effective climate in the classroom for students and faculty to pool their skills and experience. With that in mind, we selected a diverse group of students, primarily from Art and Computer Science majors, but also students from other majors who had some experience in Computer Science or Art. Students worked in pairs on a semester-long project to produce an interactive multimedia project on a topic of their choice. They were grouped according to their skills, so that ideally each team would have a complementary set of skills.

The project was taken from conception to publication on CD-ROM, including research, storyboarding, navigation, interactivity, user interface design, screen design, typography, illustration, effective visual presentation of information, mixing of media (sound, animation, stills, video), user testing, and debugging. The projects were published on CD-ROM at the end of the semester.

The collaborative nature of this project accomplished several objectives:

- The collaborative arrangement fostered an atmosphere of cooperation and communication, teaching students to work together effectively, important skills that often get overlooked in a competitive environment.

- Collaboration mirrors the process of multimedia production in the “real world.” On a practical level, this meant that students were able to fully realize a significant project within a limited amount of time.

- Finally, it set up a structure for peer learning; students helped their partners to learn the concepts that were new to them. There was a healthy give and take.

Working Models

The teams took several approaches to the division of labor. One team opted for maximum efficiency toward the goal of producing a very ambitious, complete product. These students split the work evenly between programming and design, according to the skills of the two students involved. Their project was indeed very successful, but both students admitted that they would have learned more if they had each done some of the programming and some of the art.

Another team chose the opposite approach: they split up the sections so that each student was responsible for both the programming and the art for her sections of the project. Good communication was necessary to keep the project consistent, but much of the work was done independently. The resulting project was excellent, but in this case, the benefits of the collaboration experience were concentrated in the design period. Nonetheless, this could be the method of choice for a team that does not really see eye to eye.

A third working model was best represented by the team whose credits listed one partner as the Senior Designer and Junior Programmer, and the other partner as the Senior Programmer and Junior Designer. Most of the teams took this approach, which we believe produced a greater educational benefit than the other two. Clearly, this model is the one the instructors prefer and encourage.

The first time we taught the course, we had left it up to the students to determine their own strategy. With hindsight, the second time around, we are focusing more attention on the collaboration process. The one issue that tends to arise is the equality of effort invested by the partners. We have recommended that students create a contract that can be reassessed at any time throughout the semester so that no one student feels like she is doing all the work. At the first sign of communication problems in a partnership, we suggest a meeting with the instructors to mediate.

Another issue that needs to be addressed in student collaboration projects is whether or not to permit students to choose their project partners. We had permitted such partnerships in the past in other computer science courses that were not as intensive as this one. The first time we taught this course, two of the teams were self-selected pairs of good friends. Although they did have complementary skills, we found that it did not seem to be an advantage to have friends working together in a demanding project. In fact, these two teams produced projects that were below their perceived abilities, and not as successful as the others. Our theory is that when the project needs to “go into overdrive,” friends will not push friends until it is too late. The second time

Naomi Ribner
Studio Art Department
Wellesley College
nribner@wellesley.edu

P. Takis Metaxas
Computer Science Department
Wellesley College
pmetaxas@wellesley.edu

around, the instructors formed the teams according to their perceived abilities and did not allow the students to switch partners once we made the assignments.

In an effort to facilitate peer learning and interaction early on, we structured several short assignments that required work in small groups at the beginning of the semester. At the first class session, the students were divided into groups of four and given a selection of professional-level multimedia applications to assess. They had to interact among themselves before presenting their analysis to the whole class. Exercises of this nature help to facilitate the group process for students who have been well trained for (are more attuned to) independent work in a classroom environment in which they are in direct competition with their classmates.

Throughout the semester, work-in-progress critiques are scheduled, which function according to the model of studio art critiques. It is stressed that students' participation in the role of feedback provider is just as critical as their participation in the role of presenter. The critique encourages further interaction among the students, involving them in the progress of the other projects. While the art students are well accustomed to this format, the computer science students need some adjustment to this process. This is probably due to the fact that, despite Knuth's teachings, Computer Science students often think in terms of a correct or incorrect program, and rarely in terms of programming elegance.

The critique sessions also proved to be a good time to incorporate discussion of methods of user testing in real world multimedia production. This was especially useful last year, since we did not have sufficient time at the end of the semester to adequately address user testing. This should not be a problem this year, as we have increased the length of the class sessions.

Course Contents

It should be stressed that the course was not solely focused on production. Although the students necessarily required substantial instruction in techniques and skills, there was a major focus on the theory behind the design and the programming. This theory was presented, of course, from two very different (sometimes conflicting) viewpoints, reflecting the instructors' experience and background. Interestingly, the students viewed this as one of the more notable positive characteristics of the course. We have divided the course into five major components, which are presented in an interleaved fashion.

The first component provides an overview of multimedia through case studies and introduces programming methodologies for Lingo, Macromedia Director's programming language, used throughout the course.

The second component presents the theory behind the development of hypermedia, including principles of user interfaces and visualizations of quantitative information, navigation techniques, story development, and storyboarding. We also address the appropriateness of multimedia applications for the intended purpose: for instance, when should a book be a book?

Design issues comprise the focus of the third component of the course, addressing issues of typography, design fundamentals, and color theory.

The fourth component is concerned with media technology and offers the hands-on skills and background material for working with images, sound, video, and animation.

The final component of the course considers a philosophical perspective on multimedia, touching upon the impact of technology on publishing, art, education, communication, ethics, and society in general. If there is time and student interest, the course discusses the World Wide Web and related issues.

Taking advantage of the growing activity in multimedia in the Boston area, the classwork and presentations were supplemented by field trips to multimedia research labs and local companies. Students met with multimedia professionals (digital artists, computer animators, multimedia developers) and saw the development process in action.

We have made a few changes in the structure of the course this semester as we teach it for the second time. We have doubled the contact hours of the course and organized it in a laboratory/studio format. The first 70 minutes of a two-and-a-half hour period is designated for lecture, while the remaining time is dedicated to hands-on tutorials, workshops, and interaction. We now utilize a high-tech room containing 16 top-of-the-line Macs, one per student in the course (last year, students shared computers). We also make use of other computing facilities on campus.

The Art and Science of Multimedia: An Interdisciplinary Approach to Teaching Multimedia at a Liberal Arts College

The Projects

The topics chosen by the class were very diverse in subject and in audience. We mention here selected projects implemented the first time we offered the course. A complete set can be found in the course's Web site:

www.wellesley.edu/CS/courses/CS215/frame.html

- **"Language Diversity at Wellesley College"**

The student population at Wellesley College comes from many countries around the world. This team interviewed students whose native language was other than English, videotaped them saying the phrase "I love you" in their native languages, and provided a means for the user to learn to say the phrase. Users could record their own voices and play them back for comparison with the native speakers. The students also presented writing samples in each of the languages, using a pen-pal metaphor. World maps indicated where the language is spoken, and seamless montages of background images provided a flavor of the various cultures represented. This was a beautifully designed interface with engaging original artwork. (See Figures 1-4)

- **"Souls' Midnight"**

In this children's adventure game, the premise is that twin sisters enter a haunted house. One disappears; the other (the game player) must explore the house looking for clues to help save the sister from the evil that befell her. This is another project in which the original digital paintings of each room in the house are striking. These students composed original music for their project. (See Figures 5-6)

- **"Deluxe Hotel"**

This project is about cocktail-lounge music and culture. The interface is a hotel elevator from which the user can visit four lounges. In each lounge, music of the 40s is performed, accompanied by an animation sequence. Every time the user returns to the elevator, a random selection of [elevator] music of the period is played. Graphics are beautifully done. Recipes for popular cocktails are provided, as well as information about the music.

- **"The Fractal Factory"**

This project presents the novice with an introduction to fractals. The interface uses the metaphor of a factory. The user signs in on a time card and then has the option of proceeding to the archives (where fractals are explained, accompanied by animated examples of fractals), the Observation Deck (where one can observe fractals as seen in nature), or the fractal machine. It is the latter feature that flaunts the programming talents of these students. The user enters values for the various parameters requested and then watches as the Fractal Machine creates the "custom-made" fractal on the screen in real time. While every project required problem-solving-skills in order to break down the problem into manageable, communicating pieces and implement them in independent modules, this project required more serious programming in order to calculate and display the custom fractals on demand. (see Figure 7-8)

- **"Tunnel Vision"**

A nightmarish experience in which the user explores the underground tunnel system below the college campus. Interactive animation sequences and effective use of sound effects await tunnelers as they try to find a way out. It is rigged so that no one can find a way out!

Overall, the students initially aimed very high in their project proposals and had to scale back in order to realistically accommodate the time constraints of the semester. Nonetheless, their excitement and motivation to realize their ideas was very high. Some groups set out to prove that they could indeed accomplish their original proposal in spite of our warnings, and in fact, they did.

The students worked in a small lab in the Art Department. While it was not the ideal setting, it was, a place in which they were the primary, nearly exclusive users. Working in close proximity, the students were witness to the development of their classmates' projects, and often became directly involved in the other projects during the brainstorming, troubleshooting, and critiquing sessions that happened spontaneously at all hours of the night. Excitement about the projects escalated as the deadline for completion neared. By the end of the semester, the collaborative effort had extended beyond the individual teams. Thus, the CD-burning party was a celebration of a collective accomplishment. Before burning the CD-ROM, the projects were linked through a single interface giving them a unifying theme. (see Figure 9)

Course Development

The idea for an interdepartmental multimedia course was conceived at a reception for faculty publications, when the instructors discovered, not quite by accident, that their teaching interests, projects in multimedia, and in fact, list of students, significantly overlapped. Both of us were faculty advisors for an increasing number of students proposing independent majors in Media Arts and Sciences or Multimedia Studies. These students were enrolling in relevant courses from the Art, Computer Science, Sociology, Music and Philosophy departments to piece together an independent program of study. Others were double majoring in Art and Computer Science, fulfilling the hefty requirements of both departments. Both of us had come to recognize that there was a great gap in the curriculum.

The Art Department had only one related course: a relatively new course in Electronic Imaging, in which the computer is used as a fine art tool. Some sections of two-dimensional design and photography classes used the computer for portions of the coursework. Students from these courses, excited by the technology, wanted more. Several students did animation projects as an extension of their work in Electronic Imaging, but teaching animation was beyond the scope of that course. The only options for further study of digital media within the Art Department were independent study projects or thesis projects. Furthermore, it was clear that these students would benefit from input from the computer science department.

Meanwhile, Computer Science was offering two courses aimed at different groups of students: an introductory computer science course culminating in Hypercard projects while focusing on human-computer interaction, interactivity, and functionality, and a traditional Computer Graphics course with a significant 3D design component. The Computer Science instructor felt that students could very much benefit from some guidance in the design of their projects, but it was beyond the scope of his course and the ability of the instructor.

We agreed that students in both departments needed a course in multimedia as a logical next step in their studies. We recognized that a multimedia course offered by the Computer Science department would be very different from a course of the same name offered by the Art Department. While there are merits to teaching within a single discipline, we were interested in exploring the possibilities of teaching with a multi-disciplinary approach, aiming for a richer classroom experience for the students. Thus we joined forces and applied to Wellesley College's Educational Research and Development Funding Committee for support of an experimental multidisciplinary course in multimedia.

Ideally, our course would be but one of a cluster of related courses. The students who would take the multimedia course would already have taken courses in electronic imaging and programming. There would also be courses in animation and desktop video. Unfortunately, adding courses in a tight curriculum is an interdependent process, particularly complicated when departments are required, as in our case, to keep their number of course units constant. Despite these difficulties, with the help of many excited students and with the support of the administration and the two departments involved, we managed to introduce the new course into our curriculum.

The interest among the student body is overwhelming. The number of students applying for this course grew from 60 applicants last year to 100 this year, from which we could accept only 16. Moreover, the students who took the course last year have continued to pursue their interest in multimedia both on and off campus. Most of the graduating seniors refocused their plans, and now pursue jobs and internships in multimedia. Those who are seniors this year are doing theses and independent study projects related to multimedia. Others are working on projects helping faculty to develop educational applications for their classes.

The Art and Science of Multimedia: An Interdisciplinary Approach to Teaching Multimedia at a Liberal Arts College

Conclusions

Liberal arts colleges are faced with the problem of balancing an increasing demand for new courses in rapidly evolving fields, such as multimedia, with limited resources with which to develop a new program of study. By joining forces, we were able to bring a multimedia course into existence via an experimental route. But the interdisciplinary approach of ARTS215/CS215 provides more than a quick fix to a logistical problem. The unique climate of an interdisciplinary course fosters the cross-fertilization of ideas, appropriate at liberal arts colleges with and without full-fledged multimedia programs.

The results of our experimental course far exceeded our expectations for the excellence of the projects, the motivation of the students, and the impact on the students in their subsequent studies and career paths. In the process, the students learned a great deal not only about multimedia, art, and computer science, but also an important lesson about the nature and benefits of collaboration, a subject often overlooked in today's increasingly competitive society.

The collaboration of faculty proved to be a beneficial learning experience in and of itself. It is clear that artists and programmers have very different methodologies. We have both learned a great deal from working together. A welcome by-product of the endeavor was the exchange of art and computer science ideas, which has subsequently inspired our collaboration on other multimedia projects unrelated to the course.

Figure 1 "Language Diversity at Wellesley College," opening screen

Figure 2 "Language Diversity at Wellesley College," index of languages. For each of these languages, a writing sample is provided, as well as a world map indicating where the language is spoken.

Figure 3 "Language Diversity at Wellesley College," a transliterated Russian letter.

Acknowledgements

We would like to acknowledge the students whose work we cited in this paper:

Janet Lee and Katy Ong
"Language Diversity at Wellesley College"

Colleen Baik and Eva Pedersen
"The Deluxe Hotel"

Rebecca Bargoot and Aditi Rao
"Souls' Midnight"

Lila Kanner and Alta Lee
"The Fractal Factory"

Olivia Cortina and Susan Wasseluk
"Tunnel Vision"

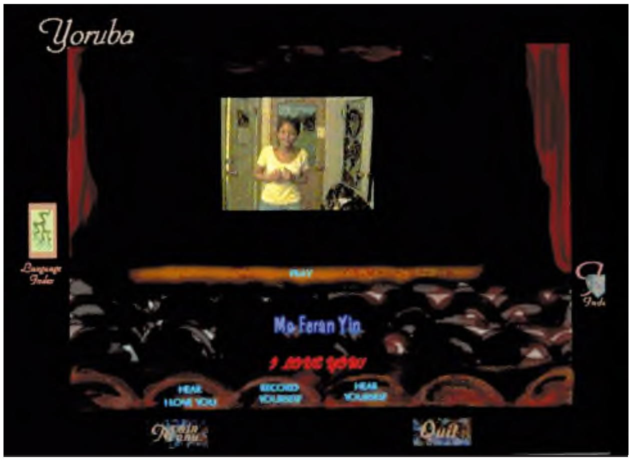

Figure 4 "Language Diversity at Wellesley College:" with the metaphor of a theater, the user is taught to say "I love you" in Yoruba, with assistance from a native speaker.

Figure 5 "Souls' Midnight," entrance hall of the haunted house, with ghosts.

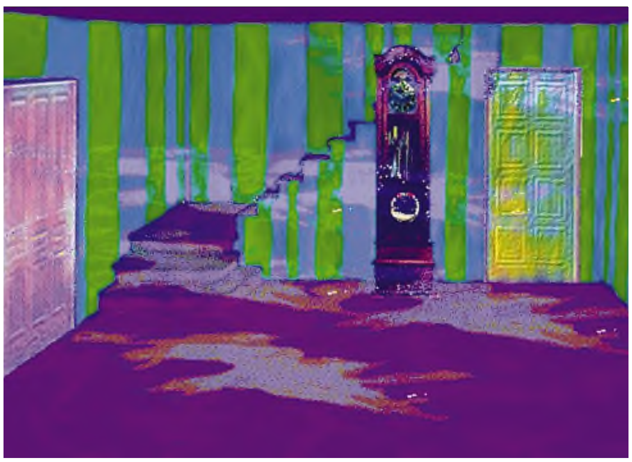

Figure 6 "Souls' Midnight," the Library. The brick is a movable sprite that was found in the haunted house. The mystery is solved when it is restored to its place in the library.

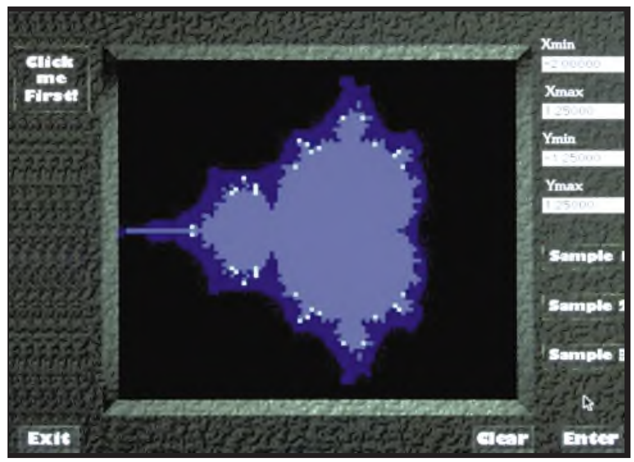

Figure 7 "The Fractal Factory:" The Fractal Machine shown here is creating a custom-made fractal.

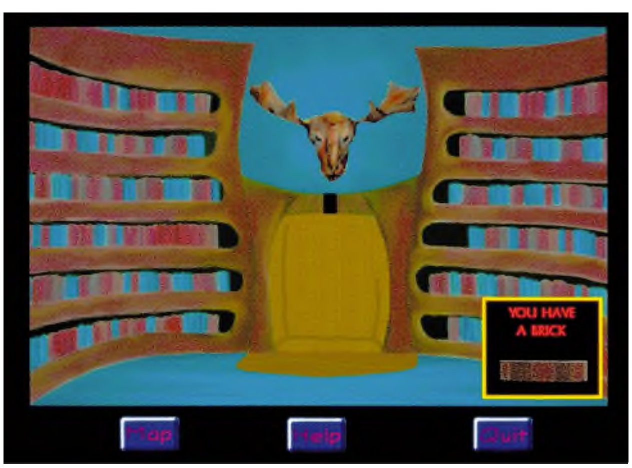

Figure 8 "The Fractal Factory," the completed fractal.

Figure 9 "The Art and Science of Multimedia," the interface of the CD-ROM from which all the projects can be accessed.

Media Technologies and an Interdisciplinary Approach to Program Design

Dave Keskeys
Department of Design
Cheltenham and Gloucester College of Higher Education, Pittville Campus
Cheltenham, UK
dkeskeys@chelt.ac.uk

"Another factor… is technical 'convergence' of computing, telecommunications, and media into an emerging digital format…"

"In from the Margins"
A contribution to the debate
on culture and development in Europe.
Published by The Council of Europe.

Introduction

In 1992, I and a number of colleagues were tasked with designing and subsequently implementing a new Bachelor of Arts with Honours degree program (BA Hons) which was to build upon the strengths of our existing provision within the areas of graphic design, photography, and video production. This new program had also to fit within the College's well established Undergraduate Modular Scheme (UMS) and was to be the first of the arts and design areas within the institution to be integrated into the modular scheme. At this time, a small group of staff, including myself, were also working with a team of architects on the design of a new Media Centre to be built to accommodate the new degree course.

This combination of circumstances presented not only a challenge, but an opportunity to re-assess the content, structure, and delivery of the media-related work offered within the institution.

Philosophy

The nature of the media industries has changed dramatically over the past decade and will continue to develop at a tremendous pace as we move into the next century. One of the primary catalysts for this rapid development has been the increasing use of and dependence upon digital technologies. This, in turn, has brought about changes in employment patterns and the skills requirements of those who work within the media. Not only should prospective media workers be confident and proficient in the use of media technologies, but they need to be versatile, adaptable, and open to a process of continuing professional development during their working lives.

Our media degree program would, therefore, aim to equip our graduates with the necessary subject-specific and generic skills and knowledge to enable them to gain meaningful employment in their chosen areas of the media. We set out to ensure that we could provide a subject-specialist practical education within an interdisciplinary learning environment and a multi-disciplinary option for students desiring more breadth to their studies. At the same time, we would aim to cater to the needs of the media industry for technically and technologically competent, creative, flexible graduates who can work independently or collaboratively as part of a creative team.

Utilisation and application of new media technologies were identified in the early stages of development as the essential components around which the program would be structured. The use of digital technologies that permeate each of the separate subject disciplines was to provide a unifying feature that would aid integration of the parts into a coherent whole. This BA (Hons) would offer graphic design, photography, and video and audio visual production with students empowered to combine courses from across these subject areas or able to concentrate on a single subject specialism.

Modularisation and Course Structure

In order to appreciate the freedoms allowed and constraints imposed upon the development team by the UMS itself, it is necessary to gain an understanding of the structure of the scheme as it operates throughout the other degree programs within the institution.

The aims of the UMS are:

- To provide a modular scheme offering students a variety of flexible programs of study in a range of disciplines.
- To offer programs of study which are intellectually challenging and which prepare students for entry into a wide range of occupations.

The units that make up an individual student's degree program are "modules" or courses with a particular practical or theoretical focus. Full-time students take 10 modules in the first year and eight modules in the second and third years. The additional two modules in the first year cover generic study and transferable skills such as, research skills, time management, and information technology skills.

The study time for each module is spread over a single semester or 15-week period such that a student would be working on five or four modules concurrently. This would produce the typical undergraduate student timetable shown in Figure 1. This pattern of delivery was not thought by the development team to be one which could easily be adapted to meet the needs of the new media degree program. The spread of the teaching and learning time over 15 weeks does not accurately reflect the short deadlines and typical working patterns and practices prevalent within many media production disciplines, which the staff wished to see emulated within the program.

We were determined to design a program that mirrored the professional practices of the areas into which our students would be moving on graduation, and we were not content to be restricted and bound by the constraints of a scheme designed to cater for predominantly theoretical disciplines.

Figure 1

Program structure 1

Level I

Semester One - weeks 1 to 15	Semester Two - weeks 16 to 30
Module 1	Module 5
Module 2	Module 6
Module 3	Module 7
Module 4	Module 8
Study Skills	Study Skills

Levels II and III

Semester One - weeks 1 to 15	Semester Two - weeks 16 to 30
Module 1	Module 5
Module 2	Module 6
Module 3	Module 7
Module 4	Module 8

The solution to the problem faced by the development team was to break the rules! A proposal was put forward to develop "short fat" modules (no sizeism intended) for delivery of the practical coursework and to run these alongside the "long thin" modules that provided the theoretical underpinning for the program. The practical modules were to be delivered in a sequential manner, allowing for gradual building of skill levels and conceptual understanding throughout the academic year and from one level to the next. The proposed structure of the program is shown in Figure 2.

The team recognised that the proposed structure (Figure 2) had one distinct disadvantage compared with the traditional UMS model (Figure 1). The timetable disparity caused by the delivery of the practical modules in five- or six-week blocks would preclude non-media-specialist students from accessing these modules. This exclusion runs contrary to the UMS objective of providing a high degree of permeability across the Scheme and facilitating access to areas other than the student's specialist disciplines:

"To encourage students to appreciate the nature of attitudes, modes of thought, and practices of disciplines other than their specialist areas."

Undergraduate Modular Scheme
Educational Objectives

Figure 2

Program structure 2

Level I

Semester One - weeks 1 to 15			Semester Two - weeks 16 to 30		
Module 1	Module 2	Module 3	Module 5	Module 6	Module 7
Module 4			Module 8		
Study Skills			Study Skills		

Level II

Semester One - weeks 1 to 15			Semester Two - weeks 16 to 30	
Module 1	Module 2	Module 3	Module 5	Module 6
Module 4			Module 7	
			Module 8	

Level III

Semester One - weeks 1 to 15			Semester Two - weeks 16 to 30	
Module 1	Module 2	Module 3	Module 6	Module 7
Modules 4 and 5 Dissertation			Module 8	

However, the development team felt strongly that the advantages of the "short fat" modules (their reflection of the working practices of much of the media industry and their intensive periods of design and production work for our students) far outweighed the lack of permeability afforded by the proposed structure. After much debate and agonising by the UMS Management Team, the proposed structure was accepted as the basis for development of an interdisciplinary media design program – in effect, a small modular scheme for media students operating within the wider UMS.

Curriculum Design and Digital Technologies - The Starting Point

In designing a curriculum to meet the needs of graduates entering an increasingly complex and changing media industry into the next century, we recognised that integration of design for and using digital media technologies should be a primary aim of the development team. Each of the disciplines to be represented within the program had previously engaged, to a greater or lesser extent, with computer technology in the content and delivery of each subject syllabus, and the staff acknowledged the inevitability of increasing engagement in the years ahead.

The graphic design program that existed prior to the development of the new degree program relied heavily upon computer technology, not only for the design and production of artwork for print media, but in the design and presentation of computer-based information resources at a time when very few programs in the UK were addressing issues of design for interactive screen-based media. A network of Macintosh computers with associated image capture and output devices formed the basis of the design and production studio. Our advertising and editorial photography program had a well-established reputation within the British Institute of Professional Photographers and UK Association of Photographers for excellence in the provision of photographic education. In the early 1990s, the photography team were just beginning to address the demand for digital imaging and manipulation within the program and were starting to utilise computers to enable their students to show their images within the context of the media for which they were intended.

The extent of the video and audio/visual work within the College was very limited, but computer-programmed tape-slide productions were being produced, and there was a growing realisation that digital editing of video and audio sequences would have to be catered to within the new development.

Clearly, a substantial investment was necessary in both the Media Centre building itself and the hardware to equip it, if the proposal was to be successful. The College determined that an innovative media provision was a worthwhile investment for the future, and the program was validated for a first student intake in September 1993, A month later, the Media Centre was officially opened by film producer Lord David Puttnam. The title of the qualification is a Bachelor of Arts with Honours in Professional Media with the specialism added for students who achieve a requisite number of subject-designated modules such as Professional Media (Graphic Design).

The Professional Media Module Map

The modules developed and now offered to students are represented in College documentation in the form of a "map" or plan of the three levels of study. All modules are available to all students registered in the program provided that they have taken and passed any prerequisite modules deemed necessary at the lower level of study. For example, any student, regardless of the specialism for which they are registered, may take a Level II video module provided that they have taken the Level I video module that provides the required knowledge and skills to underpin the more advanced work. For the purposes of this paper, information regarding prerequisites, compulsory modules, and award requirements is not included on the map. In order to simplify and maintain clarity in the presentation of the map, the Study Skills modules shown in the earlier structure diagrams and the theoretical modules (the "long thin" modules), including the dissertation, have also been omitted.

which many of the principles surrounding the content and delivery of the program are built.

All students in our media program also receive tuition in Information Technology skills within one of the Study Skills modules. This module instructs students in the use of scanners, printers, text, page-composition software, spreadsheets, databases, and research methods.

These two modules provide the core introductory knowledge and skills to enable students to profit from the more complex and involved computer design and production work and media theories covered within later modules.

Figure 3

Laying the Foundations

The Media and Technology, a theoretical module compulsory for all Professional Media students, explores the social, cultural, and practical issues around development and exploitation of digital technologies within broadcasting, the leisure industry, print media, information design, and education. The Media and Technology module lays the foundation stone in establishing the concept of the convergence of computing, communications, and media through digital technologies upon

All Professional Media students take each of the introductory modules in photography, graphic design, and video in the first semester of Level I, giving each student an insight into the working methods and practices associated with disciplines other than their own. The organisation of the curriculum and availability of module options within the timetable has been designed to enable each student to work across the range of disciplines or to follow a well-defined subject-specialist pathway in Levels II and III.

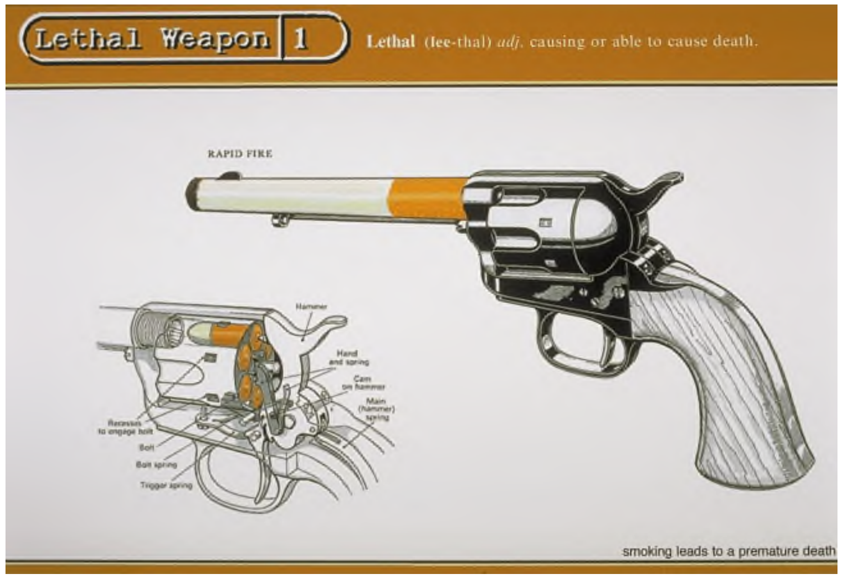

Figure 5 ‘Lethal Weapon’ by Trevor Warne (1995-98)

Practical Options

Many of the practical modules afford the opportunity for the tutor to present a variety of project briefs with the intent that students can choose the brief that most closely reflects their own interests and enables them to develop the skills needed in pursuit of their individual career goals. For example, the module Idents and Campaigns might include an opportunity for students to work collaboratively on a piece of work such as a television channel identification slot through video or computer animation, or to work independently on a public-awareness advertising campaign using any appropriate medium. Typically, the tutor might set an anti-drunk-driving campaign or an anti-smoking campaign, which could be presented in the form of designs for 48 x sheet posters or magazine page advertisements. ‘Fagtastic’ (Figure 4) and ‘Lethal Weapon’ (Figure 5) show single elements of the different solutions presented by two students for an anti-smoking campaign aimed at the teenage population. It can be seen from these examples that there is the potential for a tremendous amount of overlap in the type of work required from and produced in response to the briefs set within different modules. ‘Lethal Weapon’ might easily have been produced within the modules Image Manipulation or Advertising Agency or as an Independent Study module in Level III.

Team Working

A distinct advantage of designing an integrated media degree program over traditional stand-alone subject programs is the possibility of bringing together the talents and skills of students from different disciplines to enable each to learn from the other and to contextualise the work of each discipline in relation to the others. The development team wished to promote a teamwork ethos that echoes the multidisciplinary team approach involved in much media production work. It was therefore decided, at a very early stage in the program development, that modules would be devised that were appropriate to the work of two or all three of the major subject disciplines. Every opportunity is afforded for students to work as part of a creative team. This approach is designed to enable all students to work to their strengths, to pursue individual research interests, to produce portfolios that reflects their own career aspirations, and to place these individual aims within the context of the wider media arena. Video and tape-slide are the most obvious modules that require group work, but the program tutors endeavour to ensure that the teamwork approach is adopted in other appropriate areas of the curriculum. For example, the Advertising Agency module outline states:

“Students will be given the opportunity to form into creative teams of up to

Figure 4 ‘Fagtastic’ by Nicholas Pauley (1994-97)

Media Technologies and an Interdisciplinary Approach to Program Design

four members to work collectively towards the formulation of an advertising campaign. The module therefore encourages the interaction, in a live sense, of team members such as an Art Director, Copywriter, Photographer, Typographer etc."

Similarly, the Level III Multimedia Presentation module outline states: "Students may work as part of a design and production team within this module. Each student or group of students will research and collate the necessary textual, graphic and photographic materials needed to assemble their presentation. It is expected that each presentation will utilise the computer's facility to integrate sound, moving images, still graphic and photographic images and animation sequences in a single program."

The Use of Media Technologies

Computer-based information design and multimedia are gaining a higher profile within the program as students start to appreciate the growth in employment potential for graduates with skills in design for interactive media. To some extent, our program is demand driven, not just by the needs of the industry but by the way in which students increasingly orientate their work within the modules on offer. As our students very often have a choice in the output media they can use in response to the design problems they are set, it is easy for the tutors to assess the areas of growth and demand within the curriculum.

A wide variety of computer or screen-based projects have been undertaken within the program by students from each of the subject disciplines. In the summer term of 1997, a photography major and a video major worked together on a Web site design within an Independent Study module, resulting in the video student gaining employment in an interactive design company on graduation. A proportion of graphic design specialists choose to work almost exclusively on computer-based work as they progress through the program, experiencing design for CD-ROMs, the Internet and interactive kiosks.

Figure 6 Pittville Campus by Tim Geoghegan (1994-97)

Figure 7 Cheltenham Kiosk by Tim Geoghegan (1994-97)

Figure 8 Fire by Richard Elbaz (1994-97)

Pittville Campus (Figure 6) shows a screen shot of a prototype Web site design for our own College produced by a recent graduate, Tim Geoghegan, who is now undertaking a Masters program in multimedia design. Tim also produced, whilst on the Professional Media program, designs for an interactive tourist information kiosk for the town of Cheltenham (Figure 7), an innovative interface design for music composition, and a number of traditional and computer-animated sequences. Another recent graduate, Richard Elbaz, designed and produced a prototype of an interactive computer games magazine for publication as a CD-ROM (Figure 8), while other students have designed an interactive mail order catalogue, a virtual postcard, promotional CD-ROMs for rock bands, and interactive portfolios of their own design work. These projects have all required the individual student to become familiar with, or to work with other students who are already competent with, image capture and manipulation through software such as Adobe Photoshop, illustration software, video and audio production and capture, and basic programming using Macromedia Director or HTML. Some students have also needed to learn a 3D modeling program in order to execute their designs.

All video students and most graphic design students are now familiar with the concept of the timeline and digital storage and retrieval, either through digital video and audio editing, programming of tape-slide sequences, or computer animation and interactive work in Director. This enables each of them to be able to communicate in the same "technical language" with other students who might form members of a cross-disciplinary team.

Media Technologies and an Interdisciplinary Approach to Program Design

Image Manipulation

Initially, the Image Manipulation modules were planned to introduce photography students to digital photographic retouching and photo-montage work. However, these modules have proved so popular with graphic design students and some video students that additional runs of the modules have been provided, and the range of options available within them has been expanded. Image manipulation as a medium for illustration has become an important element of the program, with students keen to embrace the use of computer technology to produce interesting solutions to photographic illustration projects.

Drink me (Figure 9) by Katherine Hood combines traditional studio photography with computer image manipulation and is one image in a series of illustrations for “Alice in Wonderland.” Image manipulation techniques are also used within advertising campaigns, book jackets, music CD or tape covers, information design, and interactive media. Absolut Dali (Figure 10), for example, is one of a series of advertisements for Absolut Vodka produced by graphic design student Renee Le Poidevin, which demonstrate the potential of digital technologies in enabling students to successfully visualise their advertising concepts.

Wherever possible and practical, staff encourage students to contextualise the work that they do. Photography students are often required to present their images both as photographic prints and also as computer printout, showing how the images would appear with text applied as an advertisement, within a magazine layout, as a book jacket, etc. This approach to presentation of work is designed to familiarise students with the working practices of the broad range of media disciplines and to enable them to appreciate how their own discipline relates to others.

Figure 9 Drink me by Katherine Hood (1995-98)

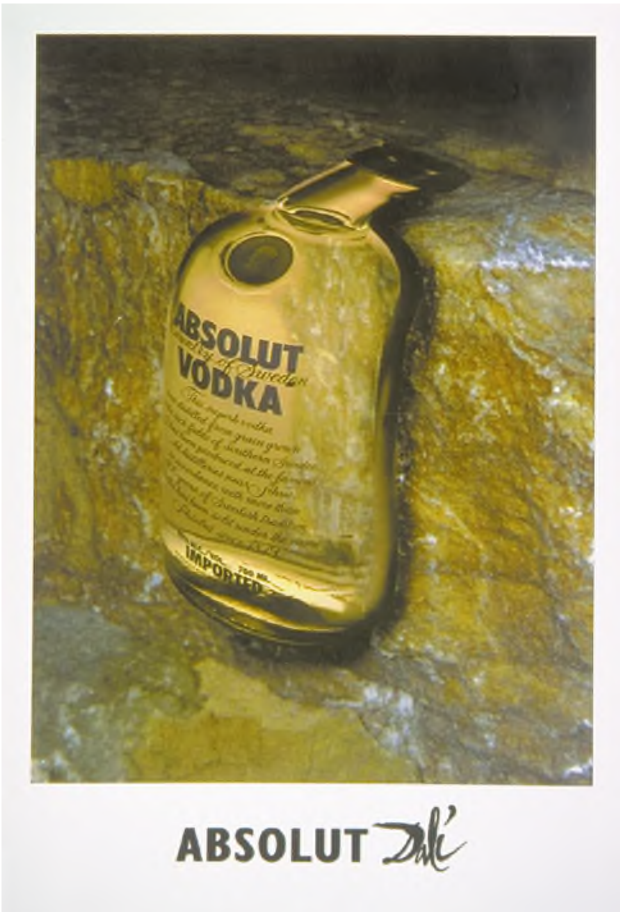

Figure 10 'Absolut Dali' by Renee Le Poidevin (1994-97)

Future Development Potential

While the initial proposal was limited to coverage of the areas of photography, graphic design, and video and audio visual production, it was recognised from the outset that the program design and structure should be capable of dynamic growth. The graphic design content of the program is currently broad-based and inclusive, with all graphics students experiencing work across a range of design areas. It was always the express intention of the development team to provide a number of more focused, in-depth specialist design options at later stages. It has become apparent that the need to provide specialist pathways in advertising design and design for interactive media are the most pressing, with an increasing number of students seeing the rapid growth of these areas as providing potential employment opportunities. The Professional Media tutors are currently working on enlarging the module base so that these additional specialisms can be offered for future intakes.

The program, as it stands, is extremely successful, having gained an excellent report from UK Government-appointed inspectors in 1997 and receiving approximately 10 applications for each place available. Our aim is to maintain or enhance the already high quality of the program as we open further options and increase participation in years to come.

Virtual Reality in Education: Irish and American Students Meet on the Virtual Frontier

Casey D. Charvet
University of Texas at Austin

J. Matthew Miller, III
Duke University

Jonah H. Peretti
Isidore Newman School

As part of a unique educational environment created by students and teachers at Newman School, in an exclusive partnership with Microsoft Corporation, Irish and American students will use multi-user virtual reality technology to facilitate trans-Atlantic meetings. This avatar-based system provides an unprecedented learning opportunity for high school history students. Historical role-playing will facilitate an exchange in which students at Isidore Newman School in New Orleans, Louisiana learn about the conflict in Northern Ireland, and students at Sutton Park High School in Dublin learn about the legacy of the American Civil War.

As society and technology change, so do educational methods. As the industrial revolution created larger, more population-dense cities, schooling changed from a largely in-home process to an in-class process. Now, with the silicon revolution, children are learning with the aid of computers. Biology books come with multimedia CD-ROMs, and digital encyclopedias outsell their traditional counterparts. As desktop computers have become more powerful, virtual reality is beginning to make its first contributions to K-12 education.

Over a year ago at Isidore Newman School, history chair Mark Cowett wanted to put the new computer facilities to good use in the history curriculum. He worked with a computer teacher who was conducting a role-playing experiment with his 10th-grade class using Black Sun's virtual worlds and multi-user technology. He had his students design alter egos, which they embodied as avatars in these virtual worlds, and report back on their experiences in cyberspace.

At Newman School, the team is now using the publicly available V-Chat software to create avatars that allow students and teachers to explore conflict resolution. In their history class, students are studying the origins of the Civil War and researching historical figures they find particularly interesting. Their research culminates as they assume the roles of Abraham Lincoln, Harriet Tubman, John Brown, Frederick Douglass, or other historical figures. Avatar-enabled historical role-playing allows students to become actively engaged in historical narratives that normally remain trapped on the pages of their textbooks.

Halfway around the world, Irish students are engaged in a parallel experience. Students from Sutton Park High School confront a different legacy. After decades of bitter conflict, Nationalists and Loyalists are still at war in Northern Ireland. Sutton Park students are becoming another cast of characters: Joseph Collins, Daniel O'Connell, and Oliver Cromwell.

Sometime this spring, an exciting meeting will occur. Sutton Park history students will enter a virtual world populated by Newman students role-playing historical figures from the Civil War. The Irish students will become historical journalists, conducting trans-Atlantic interviews with Newman's articulate and knowledgeable avatars. They will piece together complex historical narratives from the conversations they have with the student experts at Newman School. These reporters will produce newspaper-like Web pages as accounts of their virtual meetings and learning experience.

Students in both schools will discover a shared historical bond through learning about the travels of abolitionist Frederick Douglass. In 1845, he traveled to Ireland to lend support to the Irish emancipation effort. Some even referred to Douglass as the "Black O'Connell of the United States." Douglass' visit to Ireland gave him an international perspective similar to that which students at Newman and Sutton Park will gain from their contemporary trans-Atlantic journey.

Students from New Orleans will in turn interview the Irish students' historical avatars and hear first hand accounts of the Irish Civil War. This interaction will stimulate discussion of important historical themes including civil war, religious tolerance, slavery, freedom, conflict resolution, and international relations. Sharing local conflicts with a global audience will stimulate fresh perspectives on how to promote social justice.

The chief advantage of online interactive education is that multiple schools can participate. This allows conflict resolution scenarios to be explored using the same virtual world with new avatars. Furthermore, college students who have done sufficient research could supervise discussions or even assume the historical roles. The fact that the world is freely open allows for younger students to learn from the older students. Also, by gaining knowledge of different figures, students earn the right to assume that figure's avatar.

The goal of this project is to address the historical specificity of global conflict and integrate cutting-edge technology with subject-area knowledge by employing multi-user virtual environments for historical role-playing. Today we are addressing the conflict in Northern Ireland: tomorrow schools around the globe will meet in cyberspace to explore the historical roots of conflict.

Integrating Digital Technology into Classrooms:
The Making of Warp & Weft, Might & Magic, Mettle & Motherhood: An Electronic Exploration of American Women's History: 1640s to 1870s

Barbara K. Iverson
Teresa Prados-Torreia
Columbia College Chicago

Incorporation of the World Wide Web and the Internet into the fabric of public life has turned questions of whether computers and digital technology belong in classrooms into questions of how to integrate these technologies effectively to improve pedagogical practice. Our paper describes the collaboration of a women's history class and a CD-ROM multimedia production class to produce an interactive multimedia CD-ROM and a Web site, and some surprising observations about our interface metaphor and navigating in virtual space.

The project was completed at Columbia College in Chicago. As the major areas in this arts and communication college increasingly integrate computers and digital technology into their classrooms in response to the changing world of professional work, general education classes are still taught as lecture classes in "technology-free" classrooms. Our model attempts to change this in profound and lasting ways that can be easily accomplished within the usual constraints of scheduling, limited facilities, and the need for faculty development.

The Model

Our strategy is to pair up a traditional "content-rich" lecture/discussion class with a "technology-rich" computer production class to create a collaborative interactive multimedia project around the subject matter of the traditional class. For our first iteration, the classes were Women's History (Colonial to 1877) and a CD-ROM multimedia production class. A multimedia approach introduces an "almost three-dimensional" character to research-based activities, like the traditional term paper, whose future is being called into doubt (Evans, C.T. & Brown, R., p. 18.). Visual imagery, icons, and integration of meaningful graphics, sounds, and video files with text support student understanding of the textual materials and engage students in the work needed to complete the project.

By focusing on acquisition of content (the historical data) and strategic knowledge (the "how to" of interactivity and multimedia), higher cognitive levels of thinking are required of students as they construct an interpretation of the data that will communicate their ideas to their audience (Smith & Reiser, 1997.)

The Process

The "technology-rich" class, which is open to juniors and seniors who have had three or more computer classes and experience using Adobe Photoshop, Macromedia Director, was modified slightly to incorporate the "content-rich" class as "clients" who would provide the subject-matter for our collaborative, team-based production during the semester as shown below:

Course Goals & Objectives

Tech class (8-10 students in a networked computer lab with Internet access):

1 Develop leadership and cooperative skills.
2 Apply programming, graphic, organizational, and analytical skills in a production environment where time, talent, and task interact.
3 Learn to translate ideas about interface and navigation design into digital terms.
4 Produce a work series that shows engagement with principles of content and form presented in this class.
5 Make substantive, documentable contributions to group products.
6 Learn to estimate a production budget for similar projects.
7 See first-hand how a project is begun, developed, and brought to production in a real-time way.
8 Be able to map the navigations & branches of our project.

Women's History Class (25 students in a classroom with a chalkboard and desks; access to general-purpose computer labs):

1 Students are assigned to write four two-page papers, which are to include non-textual components. Teams of five students working on similar research subjects are formed.
2 Teams work together to combine their individual research papers into a single document that they then mark up with "hyperlinks" in the case of words or phrases that need a definition or other text annotation, and they identify stories or important points that can be illustrated with an image, map, quote, song, etc.
3 Enhance students' computer skills by providing a few sessions of computer orientation in the lab.

Hardware & Software

In the classroom:

1 Power Computing Macintosh (32 MB RAM/2G ROM) with internal zip drives
2 Network storage for students and the class
3 Access to Netscape & Internet Explorer browsers (WWW, email)
4 Scanner
5 Networked printer

Software:

1 MS Office
2 Adobe PageMaker
3 Adobe Photoshop
4 Adobe Premier
5 Macromedia Director
6 Macromedia Authorware
7 SoundEdit
8 Fractal Painter
9 HTML editors (Communicator, PageMill)

The paired classes did not meet concurrently, but had an overlap between their meeting times. They met together four times during the 15-week semester, and the teachers made several additional visits to their partner's class.

Teaching assistants, students from the respective classes, had no special training, but were chosen by the faculty members based on their skills and interests. The TAs answered basic word processing questions for students and oversaw the exchange of research papers, which became somewhat confusing because history students could re-write their papers, allowing "versionitis" to creep into the project. In the future, we will use a document tracking system to prevent confusion.

Using email and a Web site to communicate between class meetings, and having network storage space for shared files, was essential to the project's success. The tech class was mapping images to ideas while the idea space was being written by the student content providers via their text-based research. Navigable interactive data exhibits (branches) were generated by this collaboration.

This was a risk-taking experience for the faculty, because the educational outcomes of the process had to be negotiated among the students, faculty, and classes and determined by the dynamic established between the students and teachers. The product was an unknown quantity as well, though it was our view that process was even more important than product. Additionally, interactive multimedia is not well understood in the academic community beyond the surface level, so one's work can be easily misunderstood. Students create wonderful images, but would they be historically accurate reflections of the content being developed by the history class? This came up, but as the tech students realized that their work must please the "clients," they read the research papers and began to fit their imagery to the ideas.

One of the tech students created a graphic interface for the colonial branch of the program. The class critiqued it and approved of it. However, when our "clients" came to view it, the history teacher saw the historical anomaly right away: a woman of the 1840s plains culture stuck in a collage about the colonial years. We modified the image; the experience was valuable (Figure 1).

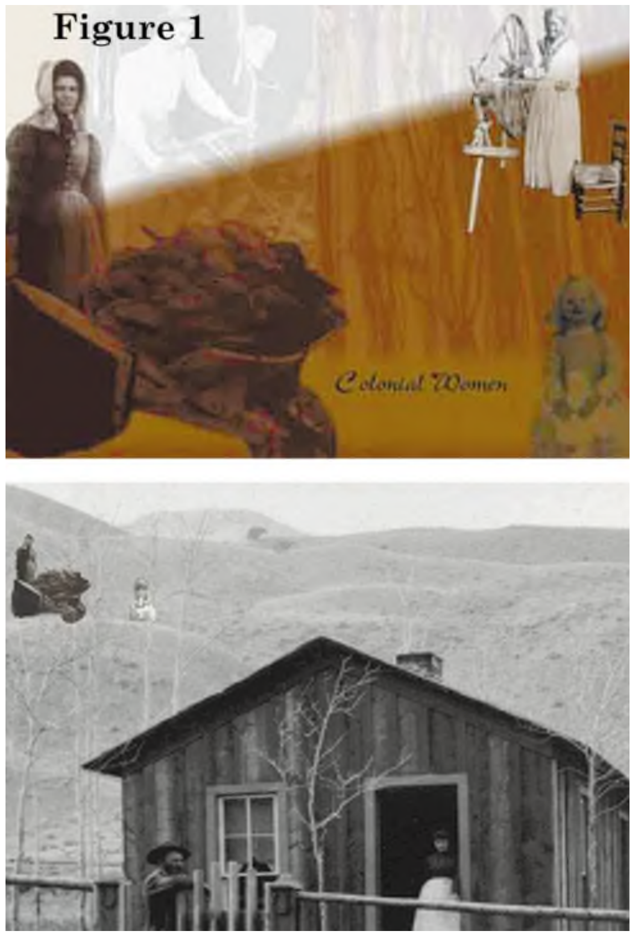

The constructivist orientation of the classes stresses teamwork and knowledge creation based on facts and other data developed by the students themselves. Traditional term papers are one-to-one communications – the student writes and the teacher reads. When students make their own interactive multimedia it implies an audience or "other." Creating a project for an audience provides opportunities for perspective-taking on the part of the student artist/researchers. This develops higher-order reasoning skills (Iverson, 1997). Research and technology skills transfer to other subjects and classes that students take. This kind of "learning by doing" helps students create cognitive webs or nets that lead to deep and full understanding of a subject (Smith & Reiser, 1997; Woltz & Palme, 1997). Furthermore, students get a sense of pride and ownership about their abilities as scholars and their mastery of the content.

We observed that history students' writing increased in length and complexity over the semester, while the tech students learned history in order to produce meaningful graphics and interactivity. Furthermore, tech students developed problem-solving skills because they were learning computer applications to use as tools for organizing ideas and content. They were motivated to get something done, and did not wait for step-by-step instruction from a teacher. Students used the teacher as a resource, but also learned to communicate with each other about technical matters. Learning became student-centered, not teacher-dictated.

The students’ ambitions on this project were limited in the main by the time constraints of a semester, rather than the constraints posed by a structured syllabus.

Students completed a beta version of the project by the end of the semester. However, several of the students tested the beta CD-ROM, and within four weeks of the semester’s end a completed version was ready for professional mastering. The professional mastering was done because we need several copies of the CD-ROM for each of the students and teachers who participated (about 35 individuals), and that is too many to do as a series of one-offs. Professional mastering also gives the students a sense of closure and accomplishment, and will provide a lasting resource for history classes at Columbia College.

Choosing a design metaphor for our project brought us interesting information about the general level of understanding of navigation of virtual spaces and the importance of paying attention to one’s audience. The tech class had to decide who we were designing our work for. We concluded that our audience was typical high-school-to-college-age students who would “have” to sign up for history. We were not making this a professional research tool, nor a piece that was aimed at a general audience as an encyclopedia might be. To attract the interest of our user, who we presumed to have an initially low interest in the topic (and parenthetically to interest some of the production students who themselves fell into this category) we developed a 3D fly-through of an historical “town” where each of its buildings would represent an era and one of our content branches (see Figure 2). Research suggests that play is a human adaptive behavior which encourages extended practice and experimentation that enables learning (Iverson, 1982.).

Navigation of the town provided users with directional arrow buttons and an

Figure 2 Interface & Navigation Details

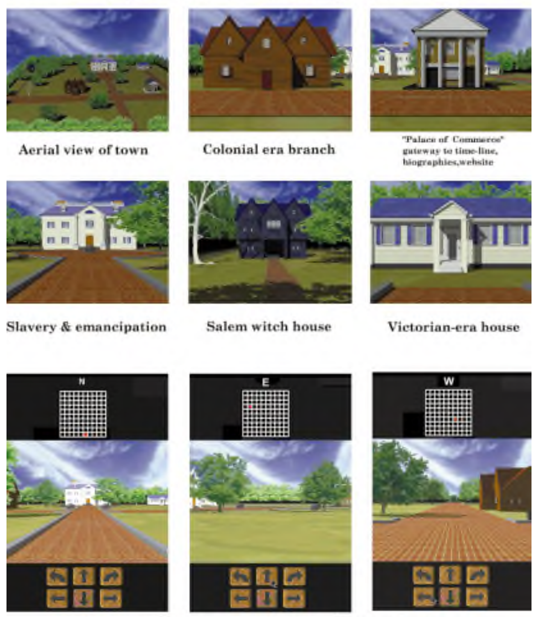

Navigation in several virtual directions

orientation indicator of their locations (Figure 2). There was no way to “jump out” of the navigation and see how many houses there were. It was up to the user to explore. Little historical context was given at this point, as our users would be expected to have little knowledge. We aimed to pique their interest with the “game-like” front-end of the program, and then once they were in a branch of the program, provide them with more content. In this way, users construct an understanding of the virtual space they are navigating. This can be frustrating, but can also encourage users to spend time in an effort to make sense of the virtual space. We felt that the more time our users spent in our “town,” exploring its branches, the more history they would learn.

This interface worked for the target users, unless they were complete computer novices. Novices were slow to have any interaction with the computer, and the interface requires interaction to “learn” about it. An un-expected problem arose from this design however. When we invited other faculty and administrators to try out the piece, we discovered a generational gap in understanding navigation of virtual spaces. The faculty (mainly over 30, and not accustomed to electronic gaming) found the navigation difficult. In fact, some believed it to be faulty. The houses weren’t labeled, they “couldn’t go anywhere” in the interface. They wanted to move directly to the content. There was confusion among these users on how the map-locator grid related to the navigational arrows (Figure 2). The faculty were not willing to navigate the town to see how many houses there were and had trouble accepting the ambiguity of not knowing exactly where they were going. They said, “How will I know if I have found all the houses?” They craved linearity. Younger users explore virtual space and more naturally construct cognitive maps of the virtual space. It is a second-nature skill to the younger users. This “sense of virtual place” is not developed, or not comfortable for older or less experienced users.

I believe the ability to move around in this interface depends on practice and experience in navigating virtual worlds. I tried this interface out with some 12-year-olds, who critiqued it in terms of its having "invisible walls," which are boundary areas of the virtual space where the programmer has not provided any visual feedback that a boundary exists. I had never heard of "invisible walls," but it is obvious that in any virtual space, there would be boundaries and that the programmer must "do something" at those boundaries (we provided a warning "beep" but no visual clue like a wall or warning sign). However, for the avid virtual "trekkers," this condition had been observed often enough that it had a name and was well-known to them.

This source of confusion between the design aims, audience experience level, and members of the academic community was unanticipated, but extremely interesting. It was of concern to me as a faculty member because my work with my students might be viewed as "faulty" instead of "targeted to a specific user community." We are going to provide a button on the interface navigation screen where those who are not cognitively ready for unaided virtual navigation can find an explanation of each "house" and will be able to tell how many houses there are. I expect that if I could collect "cookies" on who uses that button, it would tell the story of generations.

Conclusions

Most commercially produced CD-ROM software targeted to a mass audience costs from $100,000 to $300,000 per title and requires sales of millions of copies to make money. These costs discourage educators, artists, or community groups who seek to use "new media" as a form of communication. The commercial mode of production divides the viewer and producer, rendering the producer "invisible" and the viewer a "receiver" rather than a "participant." The viewer "interacts" with information that has been generated and organized by others, including information that might have been omitted. The user does little more than "look" at the piece. Knowledge is given to, not constructed by, the user.

However, when students participate in constructing interactive multimedia, it is a powerful educational experience. It yields information to students as well as experience in a collaborative team effort because it requires students to: a) generate and analyze information (text, images, sound, combinations of these), b) use meta-cognitive thinking to design the topic for interactive presentation, and c) take the audience perspective to design links, branches, and connections.

Given the widespread availability of the Internet and Web browsers that make interconnectivity between computers simple, addition of interactive multimedia presentations to course requirements at all levels of education becomes a natural enhancement of good teaching methods (Evans, C.T. & Brown, R.). The diffusion of computer skills that results from using the digital technologies as tools of production rather than as ends in themselves benefits everyone in a school. For our students to realize the mission of the college, to be the "authors of the culture of their time," they must be able to communicate effectively in the lingua franca. The pedagogy of new media is student-centered, user-friendly, and a must for our future.

Teachers have always had to teach content as well as "mechanics" (grammar, note-taking, proper attribution of source material, spelling, etc.) in any class. When computer and digital technologies are considered to be computer-mediated communication tools, and interactive multimedia is seen not as something new, but as a recent development along a continuum that began with cave painting and includes the illuminated manuscript, Guttenberg's Bible, radio, television, and film, then integration of technology into classrooms as described in this paper can be put in perspective. It is easier for students to use this technology than teachers expect (Evans, C.T. & Brown, R.), though our experience with the interface metaphor suggests that getting teachers to understand new media may be harder than we expect. More research in this area is needed.

Future Directions

We plan to continue our collaboration. We are going to provide more specific instruction in using the Web as a research tool for both classes. Content class students will synthesize their work, including identifying hyperlinks and possible images and quotes that could be recorded as sound bites prior to submitting the work to the tech class. Computer students will concentrate on some programming tasks and an analysis of the work we completed in prior semesters for several weeks before they begin working with the history content, or generating images for it. This will assure that the history students have been able to provide enough content for the work to begin. This was a problem in the first iteration, where we started off on production without having the content firmly in place.

We would like to extend this collaboration to more teachers and departments. We are working with our college's technology committee in order to secure more general purpose computer labs, as well as multimedia, presentation-ready class rooms. Our Web site will be extended with new content and improved as time goes on.

References

Evans, C. T. & Brown, R. Teaching the history survey course using multimedia techniques. In AHA Perspectives, Feb. 1998.

Iverson, B. K. Praxis What You Teach: Making Interactive Multimedia around Diversity Issues & Topics. in Gardaphe, F. (Ed.) 1997.

Iverson, B. K. Haha, Aha, Ah: A model for appraising play in curricular inquiry and evaluation. Center on Evaluation, Development, and Research, I982, No. 15, (2), 3-4.

Smith, B.K., & Reiser, B. J. (1997). What should a wildebeest say? Interactive nature films for high school classrooms. In Proceedings ACM Multimedia 1997, (pp. 193-201). The Association for Computing Machinery, New York: NY, 1997.

Wolz, U., Palme, J., Anderson, P., Chen, Z., Dunne, J., Karlsson, G, Laribi, A., Mannikko, S., Spielvogel, R, and Walker, H. (1997) Computer-mediated communication in collaborative educational settings. (pp. 51-69) SIGCUE Outlook. 25, (4), 51-69.

CAROL: Students Working on Real-World Projects Empowering Local Cultural Non-Profits

Gordon Goodman
Stephen Jacobs
Rochester Institute of Technology

In the Rochester Institute of Technology graduate course "Topics in Interactive Multimedia," students build Web sites for CAROL (Culture and Arts of Rochester Online). Sponsored by the local arts and cultural council, CAROL currently showcases the sites of over 20 local arts and cultural institutions and includes an online events calendar and a database of local artists.

Several factors make the Topics course a unique experience for students, faculty, and the organizations that participate. The CAROL consortium and Web sites are an outgrowth of the class. The course draws students from disciplines and departments within the university to work together as a cross-disciplinary team that works with representatives of the arts organizations whose sites are being developed. The representatives attend the class as both clients and students to support development of their sites and to learn to maintain and manage them after the class is over. Community investment in CAROL extends to a local service provider who donates shell accounts and Web hosting to the non-profits free of charge.

A pro bono local arts marketing campaign will use www.carol.org as its focus. The campaign will include not only Web marketing (such as banners and the like) but newspaper ads, billboards, and other traditional advertising media.

Real-world learning requires real-world problems. The Topics course is less about implementing "cutting edge" technology and more about building a site that matches the needs of the clients and their audiences. The course content, while it generally focuses on advanced Web building tools and techniques, also changes to meet the needs of the clients. The subtext of the course is collaboration. Building Web sites requires not only technical and artistic proficiency but content. Our students have the skills, and the cultural and artistic organizations have the content. The blend of providing a service for these non-profits while creating real Web sites fires the student's enthusiasm; they end up pouring their hearts into the project. Everyone wins.

As instructors, we're able to provide our students a top-flight, exciting, and engaging learning experience. At the end of the quarter, the students have significant portfolio pieces that demonstrate not only technical proficiency but teamwork and client relationship skills as well. The cultural organizations end up with not only a new and/or improved presence on the Web, but with staff members increasingly able to carry on the work and continue the institution's growth onto the net.

This paper summarizes:

- The history of our students' projects and their evolution from building individual sites to becoming the center of an organization that supports the efforts of local organizations on the Web.
- The process of running a course like this.
- The lessons we've learned in working with arts and cultural organizations and with students from different departments, disciplines, and cultures.
- Future plans.
- Some suggestions for implementation and adaptation of our model.

The History of Topics and CAROL

The history of the CAROL project is a description of process. It illustrates both conscious decisions and fortuitous circumstances that enabled the CAROL consortiums to come into being and flourish.

First Steps: Individual Museum Sites

In 1994, the Web was a fairly new and novel place. RIT did not have a campus Web server, and faculty did not have browsers installed on their computers. As he taught students to model virtual museums using hypermedia, Professor Goodman found that his assignments succeeded because they required real content and a touchstone for evaluating success. The course also provided technical challenges in terms of design, presentation, scale, and communication that transcend the individual screen. But the fruits of the class efforts were rarely seen outside of academic presentations. Moving the paradigm to the Web seemed a natural extension of the class efforts.

The first Web site experiment was done with the Genessee Country Museum in conjunction with one of their staffers, using Professor Goodman's desktop Macintosh as a server. While the site never really made it onto the Web full time, due to multiple factors including a lack of support at the administrative level of the museum and a departmental move to laptop computers, it was still a success. The success came from seeing that the goals of the project could be achieved in one quarter and that the students had an exhilarating and effective learning experience.

The next time Topics was offered, the project was a Web site for The George Eastman House International Museum of Film and Photography. This was a large-scale undertaking shared by Topics, taught by Professor Goodman, and a second course Project in Interactive Multimedia, taught by Stephen Jacobs. The institutional interest in the site's development came from Roger Bruce, the Museum's newly hired Director of Education Services and a documentary producer and media artist with a long history of NEA and NYSCA panel service. This made for a much more conducive relationship between our project and the administration in the beginning, but one that still provided some problems that we'll address later on.

Once the Eastman House site was online, it received positive press from most of the major photography publications. The Discovery Channel put a link on its Web site to ours, throwing RIT's Information Systems organization's into a panic, as they envisioned Web server gridlock. and our campus connection to the Internet was flooded by outside requests. As a result. we sought help from our community. The CAROL domain was moved to a local ISP, Service Tech (now Verio New York), which generously hosts the project, donating gigabytes of storage and considerable bandwidth.

CAROL first steps: "If you build it, they will come."

The success of the Eastman House site, and another casual remark by one of our museum clients that "a consortium to continue and expand this kind of stuff" would be a good idea sent us down the road to evolving CAROL. After a summer spent calling, and meeting with, the administration of large and small cultural organizations, two new organizations signed on to work with the Topics students. That fall, four organizations began working with the class. The Strong Museum (a fairly large museum in Rochester) and Writers and Books (a smaller literary arts center) joined the fold. The Eastman House came back for more, and the Genessee Country Museum's original site materials were resurrected, updated and rebuilt. In addition, a top-level page serving as an entry way to the other four sites announced itself as the beginning of a new consortium (the four participating organizations agreed to support this). A newspaper article on the new consortium prompted a call from the Arts and Cultural Council of Greater Rochester inviting us to join them for one of the monthly meetings of the cultural CEO's to discuss the new consortium. That meeting led to us to where CAROL is today, a collaborative organization that fits under the umbrella of the Arts and Cultural Council of Greater Rochester. The group meets monthly in the offices of the local PBS affiliate and counts the Vice President of the affiliate, the City of Rochester's Web master, curatorial and staff members from other local museums, the Rochester Philharmonic Orchestra (RPO), and others amongst its membership.

CAROL membership isn't restricted to the large organizations. Small organizations like Writers and Books, BOA Editions (a small press publishing poetry), the Visual Studies Workshop and others are also CAROL members. Some CAROL members have been our "clients," in that our students built their sites. Others, like the RPO , the PBS affiliate and the City of Rochester link to the CAROL page and collaborate on growing the domain in general. In addition to the marketing campaign, database, and calendar mentioned above, other collaborative projects, like an Art and Culture membership pass to all of the Arts Council's member organizations have gained greater momentum through the existence of CAROL.

How a Class Works: Managed Anarchy

In general, before every class begins, the professor teaching the class that quarter identifies the clients he and the class will be working with. In the earlier days of the class, we worked with the larger, more established organizations in order to get the ball rolling. At this point, we try to work with the smaller organizations that need the help more. These organizations are picked from the membership of the Arts and Cultural Council of greater Rochester. Generally, two or three new organizations are picked for a class, with additional work occasionally being done on an existing site of a previous client.

To participate, organizations must make a few commitments. First, they must agree to send a member of the organization to join the class as both client representative and fellow student. Organizations must also commit to keeping their own site current after the class is over. By following these rules, they become fully functioning members of the CAROL community. They take partial responsibility for the creation of their site, learn the skills required to maintain and upgrade their sites, and provide CAROL with another active, growing member of the community.

In the early days of the class, regular students and the organization members learned basic Web skills side by side, as part of the class. Now that our regular students are more experienced when they join the course, we often run quick two-day workshops for the outside members before the course begins to teach them the basics.

Students are divided into teams of 6-10 each. We are truly lucky in that we can open this course across departments at RIT. Normally, this brings us a mix of our own Information Technology students and Computer Science students for technical skills with students from the College of Imaging Arts and Sciences pursuing degrees in Computer Graphics Design, Industrial Design, Computer Animation, Film and Video, and Photography. In rare instances, we've also had students from Instructional Design and other fields in the class. While we don't always have this broad spectrum of skills to choose from, we can often assemble teams with a composition similar to those in industry.

As much as we can, we let student teams select the projects they'd like to work on while we balance the mix of skills in each team. The outside member becomes an integral part of the team, obtaining permissions for content and final sign-off on the site.

Course Content

The content of the course (which runs four hours, once a week for 11 weeks) generally covers Web design and implementation, group discussions, design critiques, and team meetings. Issues of tools and techniques are steered, in part, to the client needs for the quarter's designs. The fluid nature of Web technology, and the different needs of different clients, ensure that much of the course is shaped "on-the-fly" during the academic period (not for the faint of heart).

In the early part of the course, students implement strategies and techniques on their own pages, as a "dry-run" for the client pages. At the same time, they're working with the clients through the design and content collection phases of the client Web site implementations. As the course progresses, the focus shifts to the client site implementation. In the end, students are evaluated on their early work on their own sites and on the final group project. They also conduct peer reviews to evaluate each other's performance within the team. These reviews are a significant part of the grade, reducing, but not always eliminating, instances of individual load shirking.

This emphasis on real-world assignments, short production times, cross-disciplinary teams, and real-time client contact and education has proven very beneficial to the students. They end up with portfolio pieces that speak not only of their individual skills and creativity, but also their ability to succeed in an environment very much like the one they'll experience at work. It is not uncommon for students to continue to work with an organization on their site after the class has concluded. Sometimes, they finish working on aspects of the site that were not completed during the quarter, but just as often, they join the organization to extend the Web site. This real world experience within the class and after has proved enormously valuable to our students. A majority of employees within Kodak's Web Marketing Team and Xerox's Technical Documentation and Computer Based Training divisions are successful survivors of this course and the certificate and graduate programs of which it is a part.

Lessons Learned

Curators have to "learn to stop worrying and love the Web." While conditions have improved somewhat since we first began, there are still plenty of curators and other administrators within arts and cultural organizations who view an announcement of an impending Web site as they would a terminal disease. They move through all five stages: denial, anger, bargaining, depression, and acceptance. Though museums are in fact public institutions that share collections through public display, the display is under their control. The collection itself, and its representation, are the identity of the organization. Placing them in a position that appears to give control to a group of college student and/or the general public raises serious concerns on their part. Their concerns often include, but are not limited to the following:

- Piracy of assets
- Misrepresentation/amateur representation of the organization
- Concern that the Web site might lessen the public's desire to visit the actual museum

The Eastman House experience was especially instructive. Even though we had one administrator enthusiastically encouraging us to "have our way" with the site, it came at a time when the institution was between directors, and a long-time, technophobic curator was in temporary charge of the museum. This curator was not alone in her concerns of the "danger of placing digital copies of photographs" from their collections on the Web, where they could be downloaded and reproduced. We were able to overcome these objections by:

- Working directly with curators and giving them design and content approval.
- Educating the museum staff on the limited reusability of the images on their pages.
- Providing the museum the "safety net" of initially releasing the page as "a project by RIT students about the Eastman House," rather than promoting it as the official site.

This provided them with the deniability they felt they needed until they were satisfied that the site was worthy of official sanction. In the end, the Eastman House reaped the benefits of the Web site, including increased educational outreach and an expanded audience in addition to the press mentioned above, and they made the site their own. Their original staff member participant, an employee of the exhibitions department, has now had 50 percent of his time dedicated to webmaster responsibilities for the Web site. Though their site has gone through several major rebuilds (some with our students and some on their own), it still retains some of the original content created by our students for the original site. The Eastman House has become one of our biggest advocates.

Next Steps

NEA and NYSCA funds have been granted to the local arts and cultural council to support the growth of the project. The funds were granted, in part, to fund a coordinator position to focus on providing technical support for non-profits as they continue to maintain their sites and acting as web-master for the CAROL domain.

As this paper is written, we are embarking on a variation of the theme for our spring quarter. We are planning to develop a hybrid CD-ROM and Web site to support the Sesquicentennial for the Seneca Falls Women's Rights conference and the National Women's Heritage Trail that was proposed by local organizations and is being supported by a bill in Congress. The domain name we've reserved is www.womensheritage.org, and the site should be active by the time this paper is published.

Creative Expression on the Digital Canvas: An Online Digital Art Class

Jeremy Sutton
jeremy@portrayals.com
www.portrayals.com

This paper discusses how I have created and taught an online digital art class. The focus of the class is creative empowerment. The goal is to encourage and facilitate students' creative expression on the digital canvas, harnessing the unique benefits of the digital medium both as an exciting, versatile, and powerful medium for painting, and as an efficient and convenient medium for worldwide communication and interaction.

I explain the background to this class, the choices I made in curriculum structure, technology, and marketing. I provide an insight into the way the class has evolved, sharing my experience implementing the course – the successes and pitfalls, the lessons learned along the way. And I provide a road map that can encourage and assist others who are interested in teaching their own online classes, or who are considering becoming online students.

An Overview

This online digital art class, a 10-week cross-platform course, teaches students:

1. To master incredible computer paint tools (MetaCreations Painter 5.0 with a pressure-sensitive stylus and graphics tablet).

2. Painting skills, applied to creating expressive portraits from the live model.

No experience with drawing, using the tablet, or using Painter is necessary.

Students are sent a ring-bound workbook that contains detailed notes and complete assignment instructions. The Internet is the principle means of interaction. Each week, students upload completed assignments into a virtual Student Gallery on the Web, where they can view each other's work. My students and I communicate with each other via email. The class discusses relevant issues, and seeks information and advice, by posting messages onto a class email list. Each week, I select extracts from my critiques and communications with individual students, and post them on the email list.

The Digital Medium

Two distinct aspects of working in the digital medium are encountered in teaching this course: students learning to paint with digital paint rather than using traditional art materials, and distance learning via electronic means versus physically attending a real-time, in-person classroom. I shall consider each of these two aspects in turn.

Digital Paint

The computer has brought to our fingertips a number of unique advantages that make learning to paint on the computer easy and fun, including:

1. Convenient access to a vast array of brushes and art materials.

2. No need for a dedicated paint studio.

3. No mess to clean up.

4. Unlimited paper and paint (no need to feel precious about materials).

5. Forgiving medium (brush strokes can be easily transformed, erased, or undone).

6. Versatile medium (in seconds, you can go from applying oil paint to watercolor to the unique Image Hose).

7. Comfortable (the vertical screen avoids the perspective distortion encountered with a horizontal sheet of paper, and the tablet avoids the shoulder, arm, and wrist fatigue encountered when you paint on a vertical canvas).

8. Healthy (no poisonous or carcinogenic chemicals and solvents to deal with).

9. High degree of control (brush behavior can be controlled with greater precision than when working with a traditional brush or piece of chalk.

10. Large range of magnification (one moment you can zoom in and paint pixel by pixel, and then the next moment zoom out and see the whole painting as small as a stamp.

11. Record, playback, and use the creative process.

12. Share artwork across the world almost instantaneously.

13. Minimal physical storage space needed for paintings.

14. Convenient reproduction of paintings.

The main disadvantages of painting in the digital medium are:

1. Lack of tactile, olfactory, and audible stimulation. In other words, the sensual aspects of working with traditional materials, the feel of the paper or canvas, the smell of the paint, the sound of the charcoal. One student taped a sheet of paper to his tablet surface to imitate the resistive feel of working on real paper.

2. Lack of whole-body motion involved in the paint process. There is a tendency when working on a small tablet surface to become sedentary.

3. Potential UV radiation exposure, muscle ache, and eye strain associated with sitting in front of a computer screen for long periods of time.

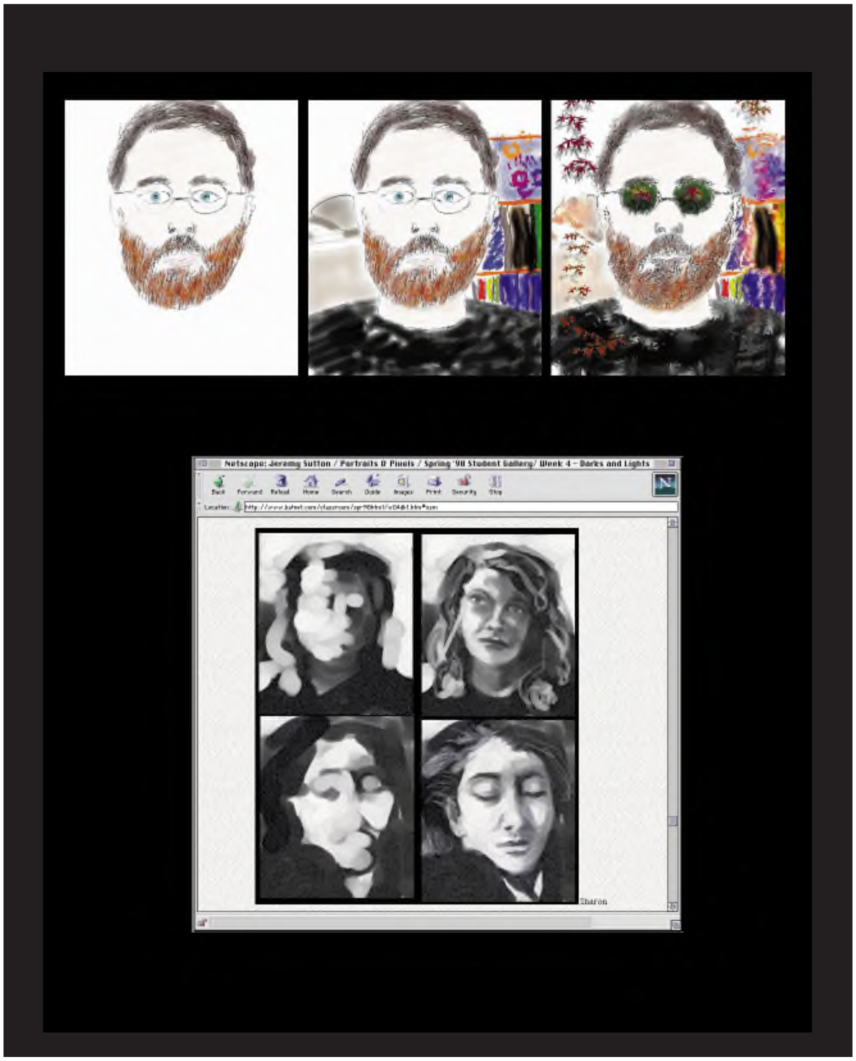

Online Instruction

Online instruction offers a number of unique advantages to students, as well as the instructor, including:

1. Geographical freedom (students and teachers can be anywhere in the world).

2. Temporal freedom (students determine their own work schedule and pace.

3. Economy (students avoid the expense, in both time and money, of going to a remote location for a residential course. The online course fee is typically less than paying for a regular class, or paying an instructor for private tuition. Internet access and email communication is relatively inexpensive.

4. Self-sufficiency-an (online course provides the motivation to become self-sufficient).

5. Interactivity at a distance (the benefit of active interactive learning and one-on-one tuition with an experienced instructor from the comfort of the home or studio).

6. One-on-one instruction at no one else's expense (typically in off-line computer classes the instructor must take time for each student at the expense of the available class time for other students).

7. In-depth critiques (the instructor provides a more detailed written critique than can be provided verbally).

8. Lack of inhibition (students avoid an instructor, or other students, looking over their shoulder as they work).

9. Simple communication (email is easy to use and immediate).

Here are student comments on why they chose an online course:

"Difficult for me to schedule my other activities around a rigid class schedule. Also the time to travel to/from class was important as I live in a fairly rural area."

"Less expensive… and I can work at my own pace without the possible concerns of how my 'work' compares to others."

"Freedom to assign my own time and place to do the lessons."

Drawbacks of online instruction include:

1. Lack of real-time, in-person contact, both between the instructor and students, and between students themselves. One student commented "I would love to meet everyone. That is the only drawback to working at home."

2. Lack of lively synchronous group discussions and critiques of each other's work.

3. Hardware/software problems (achieving trouble-free file transfer and communication via the Internet.

Creative Expression on the Digital Canvas: An Online Digital Art Class

How it All Began

In 1996, I was teaching a Painter class at San Francisco State University. During a staff meeting, there was an announcement asking for anyone interested in teaching online. I became intrigued with the possibilities. I posted a trial question on my Web site:

"I shall be teaching an online art class entitled 'Portraits & Pixels: Creative Expression on the Digital Canvas.' Please email me if you are interested in receiving more details on the class."

Over the following twelve months, I received almost 80 responses to my trial question. The magnitude of this response motivated me to make this online class a reality.

The Focus: Creativity, not Software

Prior to my online class, I had been teaching how to master tools, not how to be creative with those tools. My online curriculum focus is creativity, not software. I am launching my students on a journey in creative expression. The software and hardware tools are simply the vehicle chosen for this journey.

Guiding Principle: Simplicity

My guiding principle throughout the creation and implementation of this course has been simplicity. This has applied to the way I've structured the curriculum, organized class interactions, and harnessed technology.

Curriculum Development

The Magic of Computer Paint

My first step into creating a curriculum was to work out an organized and logical sequence of instructions that would familiarize my students with the use of the amazing digital paint tools. I needed to introduce basic concepts and guide the students through the Painter interface. My experience teaching Painter off-line helped me develop the first section of my course.

The Art of Portrait Painting

In the second section of my curriculum, the focus shifts from mastering tools to applying them. I analyzed the intuitive way I go about painting a portrait. I asked art teachers how they go about teaching their students to draw portraits. I read books to see how other artists have approached art theory and practice.

From this background research, combined with my own experience, I developed a series of practical assignments that would teach basic drawing skills, open my students' eyes to fresh ways of seeing, and lead them to create expressive and intuitive portraits.

The Workbook Concept

I considered delivering class notes and assignments via email, publishing them on my Web site, or distributing them as a PDF document on a CD-ROM. Students overwhelmingly said they wanted a hard copy workbook they could have open for easy reference while working in Painter on the computer:

"I need a hard copy to take notes in…"

"It is very important to me to have a hard copy to refer to. It is much more difficult to find something you want to refer to on the net than flipping through a manual."

The workbook is split into weekly sections, each easily located via divider tabs. The books are bound in a ring binder, which allows students to conveniently add their own notes, email printouts, and artwork. I encourage students to make notes as they go regarding favorite brushes, problems they encounter, etc. The workbook builds up into a useful resource that students can refer back to. Small graphical icons denote where there is a key point, assignment action, warning, question, or optional exercise. The workbook is easy to skim and easy to refer back to for specific information and instructions.

Workbooks are not the solution for everyone. The Corcoran Online courses have real-time lectures delivered in a virtual chat room. The University of Illinois at Chicago Biomedical Visualization online course in Electronic Illustration has all lecture notes and assignment instructions posted on their Web site.

Temporal Structure

Initially, I envisaged students beginning and completing classes at their convenience. However, as I got closer to making this class a reality I began to consider the following:

1. My need for periods of non-teaching time.
2. The complexity of instructing multiple students, all at different stages of a class.
3. The value of creating a virtual community of students, all taking the class at the same time and able to learn from each other and motivate each other.
4. The value of having distinct assignment deadlines and a defined course completion date for focusing and motivating students to complete their course work.

These factors led me to create my own semester structure. I decided to be flexible with assignment due dates, but firm with the final completion date for the semester. Corcoran Online and The University of Illinois at Chicago Biomedical Visualization online both have semesters with strict homework deadlines. The University of California Extension Online has an open-ended structure where students can work at their own pace and take up to a year to complete a course.

Workload

Instructor Workload

I am spending about 1.5 hours per student per week on direct communication. In addition. I am spending about five hours on "class business:" communicating with the whole class on general issues that have come up during the week.

My first semester, spring 1998, had a total of 15 registered students. Of these 15, three had persistent technical problems (computer crashes, difficulty obtaining software, etc.) that prevented their active participation in the course. The remaining 12 were my "active" students. Thus my total weekly workload was approximately 23 hours, and my semester total was 230 hours. This excludes the time put into writing the Student Gallery HTML code needed for each semester.

Student Workload

Most students spend between four and 20 hours per week (based on student surveys) on course study – a total course workload of between 40 and 200 hours per student. This compares with an estimated 30 student hours for a DigitalThink Photoshop 4 online course, 40 hours for a University of California, Santa Barbara Extension Painter (off-line) course, and 46 hours for a Santa Fe Photographic Workshops, Advanced Painter (off-line) course.

Students found that, in addition to taking the time to complete assignments , considerable time was needed in the first few weeks downloading software, overcoming memory problems and software bugs, implementing file transfer protocols, etc.. I have reduced the workload in the first two weeks, allowing a breathing space for overcoming technical problems.

Communication & Interaction

Commencement And Completion Rituals

I originally assumed all students would realize when the course started (it was on all the registration material). I was surprised to get an email from a student a week into one semester asking when the course was starting! This experience taught me the importance of introducing a commencement ritual.

Since that time, I always begin my courses by sending out a welcome email on the very first day of the semester. Likewise, at the end of the course I send out a thank you email.

Getting to Know Students

When students register for the course, they fill in a questionnaire that asks them about their experience, the equipment they use (important for troubleshooting), and what they hope to get out of the course. Their first assignment, a self-portrait, is the first sense I get of what they look like.

Tutorials, Guidance and Problem Solving

In my tutorials, I give positive feedback and encouragement, share my reaction to their work, comment on the assignment objectives, ask their reaction to their work, answer their questions, and make constructive suggestions. I am careful to avoid any negative judgmental statements.

An example of constructive feedback is the following suggestion to one student regarding his self-portrait assignment:

"I like your self portrait. It's well-proportioned and obviously based on careful observation. As an experiment go back to your self portrait and continue working on it: add your neck, shoulders, fill the canvas. Finally take that new version, don't look in the mirror any more, just work from your imagination. Be loose and free. Distort and transform your portrait, have fun with it!"

And his reply:

"I guess I was concentrating so hard on my head and face that it didn't really occur to me (to add neck, shoulders and background). So it was fun to go back and add to it this week."

The top three portraits in Figure 1 show the original self portrait followed by the two subsequent versions created in response to my suggestions.

Class Email List

I like to share students' discoveries and excitement with the class. Here are a couple of examples that I posted on the email list:

"I did have some trouble with the liquid metal process… I kept clicking 'OK' and then tried to apply the liquid metal, not realizing that the submenu needed to stay open…"

"Ahhhhhhhhhh… did I get surprised at myself about drawing with my left hand. At the beginning I thought my work would be unrecognizable but I discovered I have strength in my left hand and could control my arm movement to design what I saw… "

The class email list has proved a useful way for students to share problems and views amongst themselves. Here is an excerpt from an email list discussion on the difficulty of getting a live model to sit:

"Is any one having troubles getting a model? Wish I could use my cats. They are always close by."

"I'm having problems with live models too! I've set up a mirror next to my computer so that I can use myself as a model when live ones are in short supply. It was a bit of a problem with the upside-down exercise though!"

Creative Expression on the Digital Canvas: An Online Digital Art Class

"I've resorted to the cats. My cats like to curl up next to the computer while I'm working so I gave in and used them for that exercise and it worked out very well."

Group Critiques

Although opinions may be more thought out, even less inhibited, when posted to an email list, a list lacks the spontaneity and immediacy of a real time virtual chat room, or of a real-time, in-person meeting of the whole class.

I advise students who are critiquing another's work, to begin with positive statements, say what they like about a piece and why, ask about techniques they see used, and combine any other comments with constructive suggestions and solutions.

Online, as well as off-line, group critiques have a valuable role to play in helping students learn from each other, and gain confidence in presenting and explaining their ideas and feelings.

Utilizing Outside Resources

There are very useful online resources that can assist both instructors and students. Part of my role is to know where these resources are and to seek outside help when I need it. I found it useful to go to the Painter email list. There were a couple of sticky technical difficulties that responses from the Painter email list were able to resolve.

Transmitting Data

One of the toughest technical issues has been transmission of assignment file data. Most assignments involve creating graphic images (JPEG files). Initially, students sent assignment files to me as email attachments. This became impractical due to bandwidth limitations.

The solution was a dedicated ftp site that allowed students to conveniently upload their assignment files directly onto my Web server. The Student Gallery HTML documents were placed in the http directory of the class ftp site. Students now upload their assignment artwork, with carefully defined, Windows-compatible file names, directly into the Student Gallery http directory. Their work is immediately visible to me, the other students, and anyone else interested in visiting the Gallery.

Student Gallery

My original concern about introducing a Student Gallery on the Web was that publicly displaying their course work, which by it's nature is often experimental and unfinished, would inhibit students. I also did not want to introduce the complexity of a password-protected site.

The students' reaction to the Gallery has been generally positive. They appreciate being able to see other students' artwork. It's motivating for them to know how the others are getting on. It's proved instructive to see other approaches to the same exercises. One student stated: "It will make us work harder now the entire world can see what we have done." Yet another wondered about students who "are somewhat shy, not wishing to have all their work in public." She also wondered "if one is concerned on 'what to display,' several valuable exercises will never see the 'day-light.'" Some students avoid looking at what other students have done in forthcoming assignments so as not to be influenced by them.

In the Student Gallery, the artwork of each class is divided up by assignment. Within each assignment Gallery page, you can conveniently select the specific student whose work you wish to see, or you can scroll through all the work uploaded for that assignment. At the bottom of each page are all the links to all other pages in the Gallery for that class.

Students mount their own work and prepare their files for uploading. Thus they gain experience presenting their work for viewing on the Web. An example of one students' uploaded artwork is shown in the lower half of Figure 1.

Cross-Platform and Cross-Browser Issues

Software/Hardware Interaction

In my workbook notes, I attempted to cover all the obvious cross-platform differences, such as keyboard commands. I also adopted the Windows-compatible case sensitive 8.3 file naming convention. Unexpected cross-platform problems still arose.

For example, a Windows-based student brought to my attention sporadic unwanted straight lines that were appearing on images as they were painted. An inquiry posted on the Painter email list resulted in confirmation that there is a potential conflict when a Wacom tablet is used in conjunction with Painter 5.0 on a Windows-based system.

File Transfer

In one assignment, I asked students to record a script in Painter. They were to create a painting that went through transformations in time. This exercise helps the students focus on the transformative and continuous nature of the painting process. I asked them to save the script into a Painter script library file and send it to me as an email attachment. My intention was to replay the scripts in Painter on my computer and see the students' brush strokes unfold. I could not get my Mac Painter to recognize Windows Painter script files. The problem turned out to be the different ways Mac and PC create and code file information. The simplest solution was to export the Painter script as a text file.

Another aspect of cross-platform and cross-browser behavior difficulty is setting up the ftp file-transfer protocol. On the Windows platform, I recommended use of the software WS_FTP. It turns out that certain America On-line subscribers have problems using 32-bit ftp clients with Windows. They have to go into the help menu for About WS_FTP and make sure they have the 16-bit version.

The Cost

The class fee per student is $475. This includes the workbook cost, shipping within North America, an online service fee, and a tuition fee. The time spent constructing the curriculum, publishing the workbook, administering the course, and creating HTML code is not covered by the class fee.

A sampling of other online and off-line course fees (as of 3 March 1998) are:

$125.00 Digital Think (online)

$320.00 Corcoran Online (includes registration and certification fee)

$465.00 University of California Extension Online (AOL account is a prerequisite)

$710.00 University of California, Santa Barbara Extension (off-line)

$1065.00 Santa Fe Photographic Workshops (off-line)

Marketing

The primary marketing tool for my online class has been the World Wide Web. Of my first semester student body, four were alumni of Painter workshops I taught at the Santa Fe Photographic Workshops and eight students had come across my Web site and responded to my trial question.

Geographically, my first-semester student body included 11 from the U.S., one from Canada, one from the UK, one from Belgium, and one from Brazil. I have informed MetaCreations, who now list my course on their Web site. I have encouraged online and off-line reviews of my course in webzines, user groups, and high-quality international magazines.

Visitors to the Student Gallery on my Web site get a very good idea of what my course can offer. The benefit of my online marketing is that it reaches a highly targeted audience. Those who reach my Web site are already:

1. Web savvy (they have access to and interest in online interaction).

2. Interested in computer art and/or Painter (they are visiting my site typically because of a link from another computer art site or because of a recommendation from a computer graphics or Painter user group.

Thus most visitors to my site are already pre-qualified to take my class.

Evaluations

My students are my best teachers. I always learn a lot from candid student evaluations. As the course progresses, I check back with students on how it is going. At the mid-way point, they fill in a survey assessing how the course is meeting their expectations up to that point. Finally, at the end of the course,e I ask for a detailed evaluation assessing their overall impression, what they liked and didn't like, the course and workbook structure, content, delivery, assignment content and distribution, overall workload, and suggestions on how to make the course better.

Lessons Learned

The main lessons learned from my experience are:

1. Be clear about prerequisites. Warn students that they'll need dexterity in using the Web and email. Let students know all the software, hardware, connectivity, and file transfer requirements. Encourage students to make sure everything works prior to the course.

2. Keep it simple. Stick to simple technology (where possible) and simple instructions.

3. Don't assume anything! Explain everything from the basics in methodical, precise detail.

4. Structure your instructions carefully. Break every set of instructions down into easily digested steps.

5. Start with basic, simple exercises. Give your students time to acclimatize to the new learning environment.

6. Encourage interaction between students. They can be of great help and support to each other.

7. Give frequent positive reinforcement and constructive feedback.

The Future

My experience in creating and teaching my online art class has been an overwhelmingly positive one. I have been thrilled by the enthusiastic response of my students. I wholeheartedly encourage anyone to consider sharing or developing their knowledge and skills via online training.

Resources

A list of online Web-based training resources is available on the accompanying CD-ROM. Feel free to contact me, or visit my Web site, for further information on my online art classes.

Digital Image/Sound and the Fine Arts: A Double Major with Computer Science and Fine Arts

Gregory P. Garvey
Coordinator Digital Image/Sound Program
Chair, Department of Design Art
Concordia University
ggarvey@vax2.concordia.ca

This paper describes the new double major/minor in Digital Image/Sound and the Fine Arts in conjunction with the Option in Computer Science: Computer Applications. Central to the program is a two-course sequence, each six credits and two semesters in duration. DFAR 350 Multi-media Authoring was introduced in the 1997-98 academic year, and DFAR 450 Advanced Workshop: Theory and Practice in Digital Image/Sound will be offered in the 1998-99 academic year. Results of teaching DFAR 350 are presented, including a demonstration of the course Web site used to support in-class teaching and examples of student work using HTML, Macromedia Director with Shockwave, VRML and Javascript.

The emergence of the "digital arts" calls for a re-alignment and exploration of intersections across a number of disciplines previously separated by department and faculties at universities, colleges, and specialized schools. The World Wide Web is perhaps the most recent example of how a new technology "retools" pre-existing skills and aesthetic practice from previously separated domains. Successful Web page development may involve the traditional skills of a story teller, writer, poet, copy editor, art director, graphic designer, typographer, illustrator, painter, photographer, and composer, in combination with the newer expertise of the 2D and 3D animator, digital non-linear video editor, software programmer, sound effects designer, cognitive psychologist, and human factors engineer.

Howard Gardner's theory of multiple intelligences identifies seven separate modalities: linguistic, musical, logical-mathematical, spatial, bodily-kinesthetic, and intra personal [Gardner 1983]. It is clear that all modalities are used in varying degrees in multimedia design. Other useful strategies for teaching multimedia design can adopted from Donald Norman's identification of the stages of complex learning [Norman 1977] and Seymour Papert's assertion that "We should think of the computer as what you make something out of, or as a medium that gives you the opportunity to express yourself and access what other people have expressed." [Papert 1995]

Because of the diversity of programs at the graduate and undergraduate levels, the university can take a leadership role in pioneering new programs that address the educational requirements of our "digital age." The Computer Applications Option was developed in recognition that computers are being used in novel ways in widely different disciplines. The new Digital Image/Sound and the Fine Arts program was created to offer an educational "blend of computer knowledge and artistic studies." The double major program allows computer scientists to study particular fine arts disciplines, and allows artists to cross over into the more scientific realm of computer science.

This university program differs from short-term, application-specific training by offering an in-depth technical and artistic education requiring for the double major 45 credits in computer science and 45 credits in fine arts. A student can elect to receive either a Bachelor of Science or a Bachelor of Fine Arts degree. In specifying the 39-credit Computer Science Core, two credits in Technical Writing, and a further four credits in Computer Science, the Computer Applications Option provides full flexibility for students to innovatively combine a solid grounding of hardware and software with an application area of their choice.

The Computer Science Core consists of the following sequence of courses:

Year I

COMP 228 - System Hardware (3 credits)
COMP 229 - System Software (3 credits)
COMP 238 - Mathematics for Computer Science I (3 credits)
COMP 239 - Mathematics for Computer Science II (3 credits)
COMP 248 - Introduction to Programming (3 credits)
COMP 249 - Programming Methodology (3 credits)

Year II

COMP 352 - Data Structures and Algorithms (3 credits)
COMP 353 - Files and Databases (3 credits)
COMP 335 - Introduction to Theoretical Computer Science (3 credits)
COMP 346 - Operating Systems (3 credits)

Year III

COMP 326 - Computer Architecture (3 credits)
COMP 354 - Software Engineering 1 (3 credits)

A two-credit course in technical writing and one four-credit Computer Science Elective taken from courses numbered above 300 complete the 45-credit requirement in Computer Science. The remaining 45 credits can be used to complete a Major or a Minor program in any discipline of Arts and Science, or it can constitute the 45-credit Major in Digital Image/Sound and the Fine Arts. The 45-credit program in Fine Arts is designed so that students enrolled in the Major in Digital Image/Sound have a range of choices in their first year from introductory courses, which include Photography, Design Art, Interdisciplinary Studies, Computer Music, Film Animation, and Studios in Fine Arts (SFAR) courses. In the subsequent two years, students take advanced courses from the same Fine Arts programs.

Concentration requirements include the following three courses:

FFAR 250 - Visual and Performing Arts in Canada
(6 credits)

DFAR 350 - Multi-Media Authoring in the Fine Arts
(6 credits)

DFAR 450 - Advanced Workshop: Theory and Practice in Digital Image and Sound
(6 credits)

12-15 credits are taken from the following courses:

DART 200 - Design Art Theory and Practice
(6 credits)

EAMT 205 - Electroacoustics I
(6 credits)

FMAN 304 - Animation II
(9 credits)

IDYS 200 - Studio Seminar in Interdisciplinary Studies I
(6 credits)

PHOT 200 - Foundations in Photographic Vision: Theory and Practice I
(6 credits)

Or a choice from one of the Studio in Fine Arts Courses:

SFAR 250 - Visual Language as Content
(6 credits)

SFAR 260 - Integrated Drawing: Intersections
(6 credits)

SFAR 270 - Extended Studio Practice
(6 credits)

The remaining requirements include an additional 12-15 credits of upper-level courses chosen from the following programs: Design Art, Film Animation, Interdisciplinary Studies, Music, and Photography. The 30-credit minor option in Digital Image/Sound is intended for a limited number of students in Fine Arts who are not enrolled in the Computer Application Option. Similar to the major, it has two required courses (FFAR 250 and DFAR 350). The remaining 18 credits offer a choice of six to 12 credits of the same introductory courses as the major and an additional 6-12 credits of advanced courses to give students a grounding in the aesthetic and technical issues of a particular discipline.

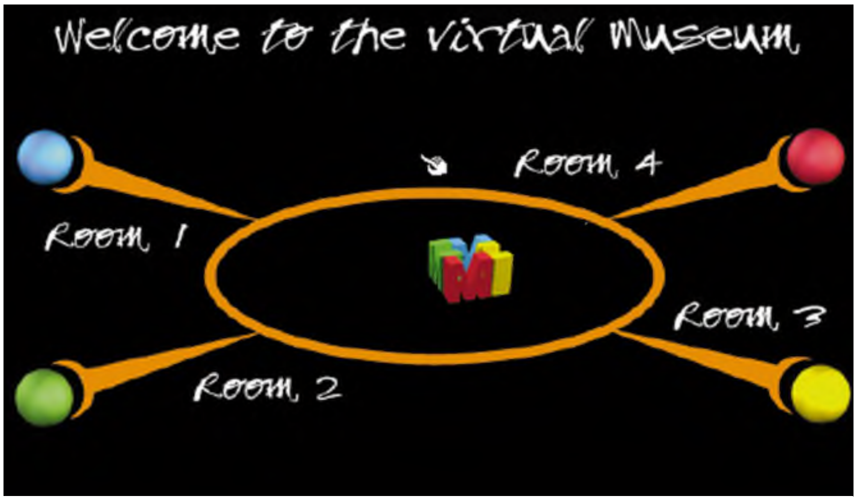

The Virtual Museum, Paul Ortchanian, DFAR 350

Admission requirements for the Computer Science Option include at least two previous courses in calculus and one in linear algebra, a G.P.A. of 3.0 and previous enrollment in provincial (Canada) preparatory programs (CEGEP 10.12) or the equivalent. For the Fine Arts, an application requires a letter of intent, transcripts, and submission of a portfolio of 10-20 slides of visual work and/or a video and/or an audio tape. Students are admitted who have a background in the fine and performing arts and demonstrate the aptitude for learning the technical skills for multimedia design.

DFAR 350 Multi-Media Authoring in the Fine Arts focuses on teaching Web page design using HTML, multimedia design using Macromedia Director with Lingo scripting and Shockwave, and building virtual worlds using Virtual Reality Modeling Language (VRML). Additional topics include issues in graphic design, color, typography and layout; perception, cognition and human factors; human-computer interface design; navigation; and hypermedia design. Sound digitizing, editing, mixing, and MIDI are also covered as an important part of multimedia design. It is expected that HTML (Hypertext Markup Language) and VRML will be learned by building Web pages and virtual worlds by hand. Other techniques are introduced, such as GIF animations, inline QuickTime, and embedded Javascripts.

Digital Image/Sound and the Fine Arts: A Double Major with Computer Science and Fine Arts

Weekly assignments are to be compiled and presented as an electronic journal and posted on the course Web site at year's end. These include exercises that develop specific knowledge and technical skills according to the current subject under study. There are two minor projects, each four weeks in length. In the first minor project (fall semester), titled "Le Musée Imaginaire," students worked with Macromedia Director to create an imaginary museum or a "museum without walls" as an interactive desktop presentation. This minor project was followed by the Fall Major Project I: "Vers un musée virtuel interactif" (Toward an interactive virtual museum). In the traditional museum setting, visitors are normally not allowed to touch the works of art. In the virtual museum, the user is encouraged to "touch" the immaterial objects. In this major project, the basic experience of looking is haptically extended by allowing the user to interact directly with and manipulate the art work.
In "Chair-not-a Chair," students selected a photograph of one of the classic "De Stijl" chair designs by Gerrit Rietveld, such as the "Blue Red" or "Zig-Zag" chairs, and modeled it in VRML, adding sound, behaviors, animations, and environments. The second major project is a student choice project, which permits students to continue developing their first major project, or continue with the VRML project, or create an entirely new project using a combination of HTML, Javascript, Director, and Lingo. Students are required to identify the project as a Web site or a stand-alone desktop presentation presumably for CD or DVD. In effect the minor projects are treated as opportunities to create in-depth experiments, mockups, and prototypes for major projects. The purpose is to build, test, and discover what works or doesn't work technically, cognitively, and aesthetically.

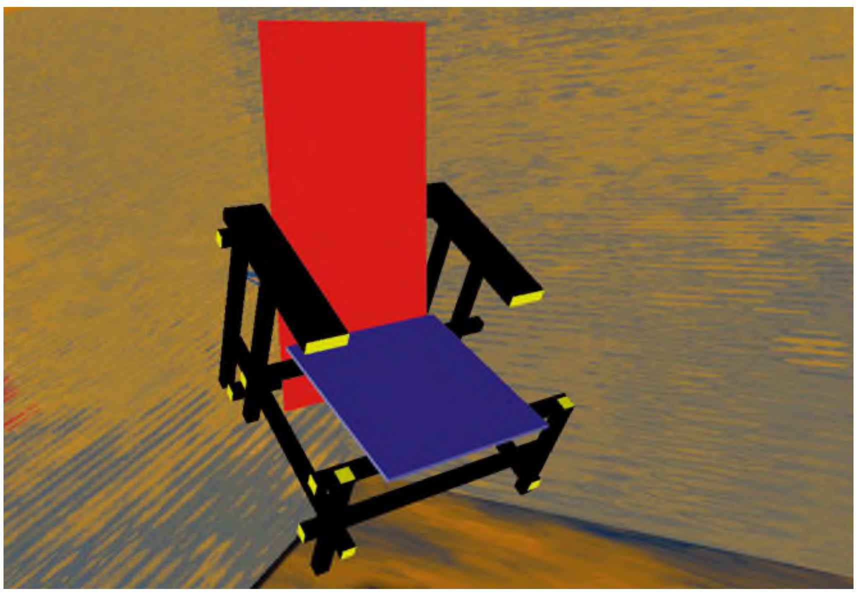

Chair-not-a Chair Yves Gigon, DFAR 350

In an special arrangement, students are given access to resources and original source materials provided by the Montréal Museum of Fine Arts, such as digital scans from the collection and authoritative text in French and English. Successful student projects may be included in the Montréal Musée des Beaux Arts Web Site. In addition, students from the Digital Image/Sound program are participating in a multimedia collaboration between Concordia University and K-12 schools sponsored by the Charles R. Bronfman Foundation and the J. Armand Bombardier Museum and Foundation. Other students enrolled in this program have been awarded contracts by Stentor Communications, an alliance of 11 Canadian telecommunication companies.

After reviewing student progress and the difficulties encountered during the first year of the program, the following recommendations for curriculum additions and changes can be identified:

Year I

DFAR 200 - Introduction to Computer Fundamentals in the Fine Arts
(3 or 6 credits)
Covers basic skills such as scanning, image processing, vector and bitmapped graphics file types, digital typography, and fundamentals of digital sound.

DFAR 250 - Introduction to Multimedia Authoring
(6 credits)
Introduction to Macromedia Director using menu-driven authoring. Web page design using Page Mill (or equivalent).

Year II

DFAR 300 – Building Virtual Worlds in VRML
(6 credits)
One year in-depth introduction to VRML with the requirement of building a major project.

DFAR 350 – Intermediate Multi-media Authoring
(6 credits)
In-depth HTML, Director with Behaviors, Lingo, Shockwave, and Lingo for the Internet.

Year III

DFAR 400 – Multi-Media Scripting and programming

(6 credits)

This course separately covers Javascript, Java, and CGI programming (Perl) for multimedia design.

DFAR 450 – Advanced Workshop: Theory and Practice in Digital Image/Sound

(6 credits)

Year-long projects using all authoring tools.

DFAR 460 – Seminar in Critical Issues in Art and Technology

(3 or 6 credits)

This courses investigates the theoretical and critical issues surrounding art and technology.

DFAR 498 – Internship Program

(3 or 6 credits)

Students volunteer at companies with faculty supervision.

DFAR 499 – Coop Program

(3 or 6 credits)

Students are employed and paid by companies to work on multimedia products.

While it is not possible to cover all the related disciplines in depth, it is clear that the multimedia artist/designer needs to be conversant with a wide range of technical and aesthetic practice. The new program in Digital Image/Sound and the Fine Arts, by offering a double major with Computer Science, provides an important step toward preparing students for a future in multimedia design. Guided by learning theories of Gardner, Norman, and Papert, this program follows the dual principles of “learning by doing” and “user-centered education” in an environment of structured collaboration. [Garvey 1995].

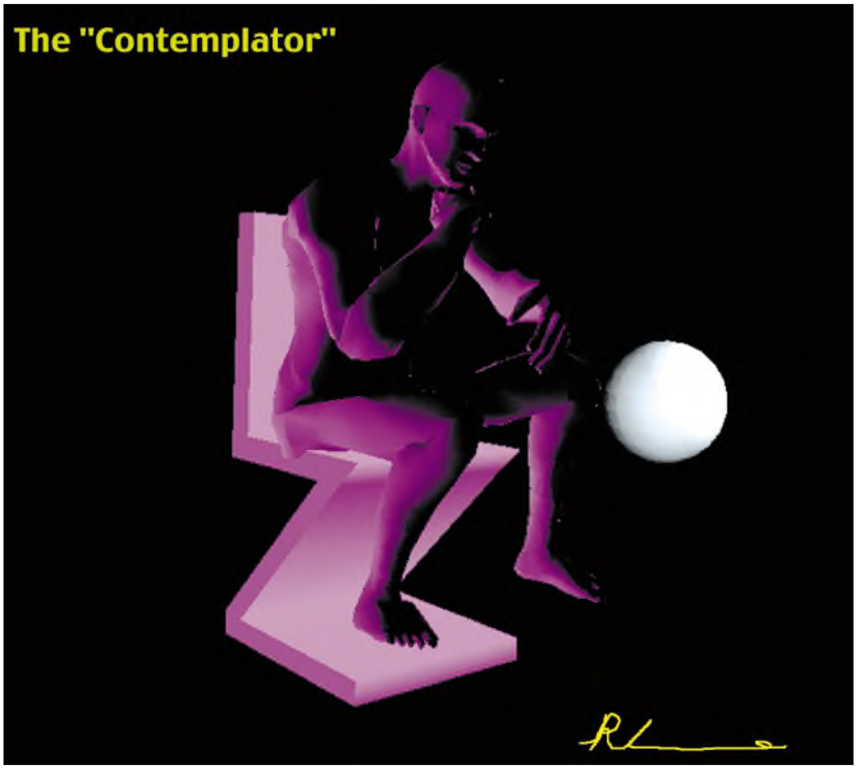

Chair-not-a Chair Roland Lowe, DFAR 350

References

Gardner, Howard (1983). Frames of Mind: The Theory of Multiple Intelligences, New York: Basic Books, Inc. Publishers.

Garvey, G. (1995) Toward a Theory Computer Design Education SIGGRAPH 95 ACM SIGGRAPH Education Committee, Copyright 1995.

Norman, Donald (April 11, 1977). Teaching learning strategies, Mimeo, University of California, San Diego.

Quoted from Multimedia Today: The Sourcebook For New Media Power Vol.III, Issue 4-1995 p.38.

The Language of Cinema and Traditional Animation in the 3D Computer Animation Classroom

Pamela Turner
Communication Arts and Design
School of the Arts
Virginia Commonwealth University

Sophisticated technology, particularly in the area of three-dimensional computer graphics, has become more affordable, enabling university art departments to add 3D computer graphics imaging and animation to their curriculum. Offering a productive learning environment where content and meaning is not overshadowed by the necessary volume of technical information can be a challenge. To use the equipment creatively, students must learn certain steps and processes. A structure is needed to simplify the complex balance between the technology, animation techniques, and content. How do students learn to use these advanced tools and at the same time develop their own visual language? The desired structure can come from identifying where traditional media, specifically film and animation, overlap 3D computer graphic animation and where 3D has "rules" of its own. From these intersections, an effective strategy for teaching 3D animation can be created.

In the curriculum where I have developed this approach, students only have two semesters of 3D coursework. The first semester, they learn to model, light, and apply materials and textures; in the second semester, they animate. More than a few are mesmerized with the technology at first, then frustrated as they try to get it to build the cool objects they had envisioned. By emphasizing their combined role of director, camera operator, lighting technician, and animator students see more of the "big picture" of moving images in 3D. The technology, the "HOW" of using the computer to create this communication, is then learned in the process of exploring film and animation techniques.

Cinema

Cinema is the art of the moving image. There is a camera that defines the viewing frame, a set for the action to occur in, objects, characters, and lighting. By looking to cinema, we can identify the basic tasks, as James Monaco does: "Three questions confront the filmmaker: what to shoot, how to shoot it, how to present the shot."[1] The "what" is tackled through conceptualizing and scripting. The "how to shoot it" addresses the use of the camera to carry the meaning and intent of the script. "How to present the shot" moves us to syntax, or how the pieces will be organized to form the complete whole. The storyboard is the template that evolves from the answers to these three basic questions. As in spoken and written language, a structure is formed that the viewer follows and feels - a beginning, a middle, and an end. This engages viewers, carrying them through conflict and resolution, in humorous or dramatic situations.

The cinematographer's tools include sets, lighting, camera angles, and camera movement. These are also the tools of the 3D animation student. Like the filmmaker, they use a script, storyboard, and concept drawings to model their scenes and characters efficiently. Scripts and storyboards are not the students' favorite thing to do! The real necessity for scripting and storyboarding is emphasized through thinking from the perspective of creating shots. To do this, the student has to consider many questions. Where is the camera (the viewer) going to be? What is the viewer's role? Will the camera move and how? What is in the scene?

From the filmmaker, we learn to plan lighting, camera angles, and movement as these shape meaning and create the desired emotion and interpretation. Color, contrast, angle, and field of view all work to give viewers information about what they're seeing and how to see and interpret it.

Translation of fill, rim, and key lights into the software's lighting tools is explored. Most 3D software has a selection of types of lights, from point lights to the sun, which the user can define. "Wattage" can be determined, as well as placement and direction, depending on the type of light. Lights can be designated to include or exclude certain objects, a task that is much harder in real life!

Figures 1a and 1b. Two examples of how lighting impacts the interpretation of a character. Images by David Beeler.

Figures 2a and 2b. This demonstrates the same set shot two different ways. In both shots, the camera starts in the same location; one shot uses tracking, in which the camera follows the object, the other uses a pan, where the camera turns on its axis to follow the object. Images by Pamela Turner.

Once modeling, texture mapping, and lighting are learned, the student moves to animation. Camera movement is the first animation lesson. Here we address the "how" of shooting the scene. Basic camera moves (dolly, tracking, pan, tilt, and zoom) and how these moves and combinations of them can be emulated within the 3D computer environment are demonstrated. There are more active camera movements, thanks to helicopters, cranes, and remote control devices, allowing the viewer to see things, to be places, they would rarely see or experience in real life. 3D software makes this even more possible. However, the new student usually has no problem making the camera fly around like an angry bee, creating a viewing experience conducive to motion sickness. Moving the camera is easy. You just draw a path and adhere the camera to the path, and off it goes. It is control and communication that is difficult to achieve. While trying not to deny students any of the shots that a free-moving camera can offer, the angry bee approach is discouraged. The emphasis is placed on why the camera is moving, not just how do we move it within the software. Is it to follow a character, to place the viewer in the middle of the action, or to reveal objects?

Students are shown snippets of film to clarify camera movement and to help them better understand how the camera is being used in the communication process. One of the examples used is from Kurosawa's "Dreams,"[2] specifically the dream of The Village of Watermills. Here the camera is an active observer, following the character as he walks into the village. The camera tracks with the character as if on a parallel path. Then the camera combines the tracking movement with a pan as the viewer's path shifts direction and converges with the character's. The viewer is then facing the character, watching as he crosses the old bridge over the stream. The camera allows us to share in the experience of the character. We, like him, are wandering, curious and enchanted by the old village.

Another example is a scene in Jim Jarmusch's "Down By Law."[3] In contrast to Kurosawa's camera, here there is minimal movement. Jack and Zack are in jail. Jack is talking about what he's going to do when he gets out. He paces front to back, left to right, defining the space. The camera only moves slightly to keep him in frame as he comes front and right. The focus is on him as Zack sits in the background. The stillness of the camera, the tightness of the frame, emphasize the relationship and isolation of the characters,

Figures 3a and 3b. Camera movement to pull the viewer into the scene as it pulls back to reveal the subject. Stills from animation by Russ Honican.

as well as the intensity of Jack's reverie. The camera is unobtrusive. The viewer is present, quiet and intently observing.

Through watching these segments, the students become more sensitive to the role of the camera, and how important camera placement and movement is in conveying meaning and intent. The camera is the viewer, and its placement and action determine the viewer's involvement in the scene. Now the software can be approached with a better understanding of the medium. How to get the camera to pan, track, and dolly. How to combine movements like pan and zoom. How to fine tune those movements, to allow the camera to hesitate, slow down, or speed up. These are learned through technical demonstrations and then by applying them through individual animation. Camera movement is keyframed, or cameras are placed on a defined path. Function curves are edited to refine the motion.

Perhaps the most difficult topic to teach is syntax. This is the structuring part of the process, where images are sequenced to communicate time and space, creating a style and structure for the work. The students have decided what to shoot and how to shoot it; now they must present these shots. Here the need for previsualization, scripting, and storyboarding becomes most apparent. They must organize their scenes and think of transitions, fades, and cuts, and plan their animation accordingly. To do this, they have to think of the animation as a finished piece and break it down into shots. Scenes must be managed so that camera movement, models, and lighting match from cut to cut. These considerations determine what sets and models actually need to be built.

This is not only a technical concern. Choices in image structure and sequence are tools, just as lighting and camera movement are tools, for communication. Fast-paced cuts indicate excitement, a building of suspense. Longer shots allow thoughtfulness and rest. A cut from a shot tracking a character to a shot from the character's point of view gives us a "second camera," access to another perspective in experiencing the scene. Addressing syntax gives the student a better grasp of structure, timing, and storytelling - elements that are usually weak in early attempts at animation, video, and film.

Students must visualize how to shoot and present the shot as they develop their storyboards, before they start modeling. They are encouraged to use the software as a tool to do preliminary tests in making these decisions. Once the storyboard is developed, it is used as a guide in the process of creating the scenes. Environments are modeled and lit much like film sets, keeping the camera in mind. Objects that appear in the background have simpler geometry and lack close-up details, as they're never viewed closely. Close-up shots are modeled separately. Walls or areas that are never seen are never modeled. Here the need for previsualization, scripting, and storyboarding becomes obvious. Otherwise, the student tends to work more sculpturally, modeling whole rooms and objects without respect to the ever-present camera frame.

Traditional animation techniques are addressed to help students create and control object/character movement. Cartoons and animation are so much a part of their visual vocabulary that they don't consciously notice things like anticipation and exaggeration. But the lack of these elements can make an animation difficult to "read." Viewers miss important signals. They need a cue that the character is going to run off screen or turn her head, a cue as to what to look at next. And even though students may have learned these techniques in traditional animation classes, they sometimes have to be reminded to apply them in 3D animation as well. These animation concepts can be difficult to simulate, depending on the student's 3D modeling ability. If the object wasn't built with movement in mind, it probably isn't going to move convincingly.

Animation videos are used to demonstrate the use of walk cycles, squash and stretch, anticipation, exaggeration. All types of animation are viewed, but I find clay animation to be particularly helpful as it, too, is 3D. A favorite example is Nick Park's "Creature Comforts,"[4] where zoo animals gesture, sigh, blink, and scratch as they are interviewed about their accommodations. Here we find wonderful examples of character motion creating personalities.

Animating in a virtual 3D environment that has no gravity is much like creating drawn or clay characters when it comes to communicating weight. The animator has to imply weight through the object's movement. Motion is studied, particularly head and eye movement, and walk cycles. The use of overlapping motion, which is critical in 3D characters, is emphasized.

Once these techniques have been identified, the appropriate software tools are introduced. There are numerous ways to create object movement, depending, of course, on the software used. Models can be translated, rotated, and scaled. They can be squished and stretched in numerous ways. Many software programs offer animation tools that create squash and stretch based on the object's keyframed movement. Natural forces such as wind and gravity can be applied so that the object reacts to the environment. Inverse kinematics[5] is taught as we investigate walk cycles. It is the students' storyboards that guide the technical demonstrations, their ideas providing the motivation to explore these digital tools, as they work to create the illusion of life.

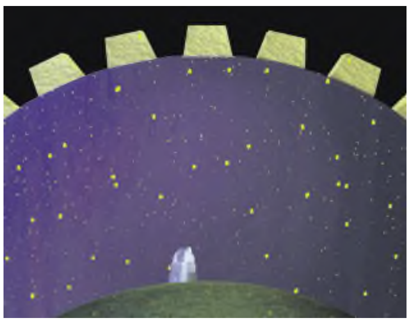

Figures 4a and 4b. Camera tracks back to reveal that things are not what they initially seemed to be. Stills from animation by Kelly Perkins.

Conclusion

Borrowing from the already-developed principles of cinema and animation has proven to be an effective approach to teaching 3D animation and more. Students are exposed to new ideas through discussion and viewing contemporary work in animation and film. This approach helps make them aware of the context they live in and gives them the tools to create within that context. Through observation and application of cinema and animation techniques, they come to understand the connection between form and content, and how these techniques are used to effectively convey intended meaning. This encourages content development, without which the work is just a technical exercise. Understanding this, they are then able to articulate complete thoughts and not simply rehash what they observe in current 3D venues (games, special effects sequences, animations, and cartoons) but to express their own individual ideas and vision.

Students' response is positive, although I know it is challenging for them as the 3D software is unlike other software they've used. The interface is different and there are many new techniques to learn. Looking at the big picture gives them a different perspective, so they can step back and view the technology critically and avoid being overly enchanted by it. I believe they become more critical of their work as they gain understanding of why and when to use this tool and not just how. This is evidenced in the critiques. There is much more discussion about the intent and communication of their work than before. I hear comments like: "Is there a reason the camera is moving there?" "Perhaps some shoulder movement would help us read her character better." "I think the camera needs to hold a few seconds longer here so we can connect to that character." They're directing and animating, not just honing software skills. I have to remind them to ask technical questions relating to the use of the software.

Using cinema and animation as a guide in teaching doesn't necessarily add to the amount of information I have to cover in the class. Rather it provides a natural structure for the techniques that are taught. Students have only two semesters total of 3D imaging - one for modeling, the second for animation – so there is not the time, or resources, for doing full feature animation here! They are reminded to integrate skills learned in their other film, animation, and video classes. They learn to model concisely, and to say something economically. Television has taught them that 10 seconds is enough time to make a point. Being able to "say something" gives them more initiative to learn how to use the technology. By the end of the semester, the computer has become almost transparent. The emphasis is on the content, not on pushing buttons.

Ultimately, I feel the students receive a better visual communications education. They get to take a further, more interesting journey. And along the way they create some pretty interesting animation.

References

1 James Monaco. How to Read a Film. (New York: Oxford UP, 1981) 148.

2 "Village of the Watermills" from "Dreams." Dir. Akira Kurosawa. Warner Bros., 1990.

3 "Down By Law." Dir. Jim Jarmusch. Perf. Tom Waits, John Lurie and Roberto Benigni. Island Pictures,, 1986.

4 "Creature Comforts." Dir. Nick Park. 20th Century Fox, 1989.

5 Inverse kinematics is an animation technique based on the creation of "skeletons" or structures that move in relation to each other, much like our bones move. This allows a 3D-modeled thigh and calf to "behave naturally" and follow the foot when it is moved.

Teaching Computer Graphics with Spreadsheets

Francis T. Marchese
Pace University
fmarchese@fsmail.pace.edu

Abstract
Spreadsheets are a great way to introduce computer graphics concepts to computer science students. Through direct manipulation of numbers, students develop a more concrete understanding of the data they compute from the formulas they derive and use. This paper presents some experiences using spreadsheets for in-class demonstrations and homework assignments.

Background
The introductory course at Pace University is offered as an elective to undergraduate students enrolled in the CSAB accredited BS and non-accredited BA degrees in computer science. Typical textbooks used in the course are *Computer Graphics* by Donald Hearn and Pauline Baker, *Computer Graphics* by Francis Hill, and *Computer Graphics: A Programming Approach* by Steve Harrington. Prerequisites for the course include completion of a computer science core including C++ programming, data structures and algorithms, computer architecture, and operating systems. The mathematics prerequisite is two semesters of calculus. Most juniors and seniors who enroll in computer graphics have completed courses in discrete mathematics and statistics. A single, two-semester laboratory science sequence such as chemistry or physics is required for the BS degree but not the BA degree. Because of scheduling requirements, the laboratory science sequence is frequently deferred until a student's senior year.

A student entering the computer graphics course with such a background exhibits one significant deficiency: inability to visualize the kinds of numbers that enter into and arise from the equations they derive and program. There are two reasons for this deficiency. First, computer science courses concentrate on symbol manipulation and processing. Numerical methods are barely considered at all. Second, the numerical data processing skills taught in chemistry and physics courses, especially through their laboratory components, have not been cultivated in these students.

Since computer graphics is part data structures and algorithms and part mathematical methods, the student deficient in numerical data processing skills is unprepared for half the computer graphics course. The programming process compounds this problem, with its design, code, debug, and test cycle, by short circuiting a student's attempt to gain a feel for the mathematical methods and the numbers arising from them. To solve this problem, spreadsheets have been integrated into lectures and homework assignments, covering topics such as: absolute and relative addressing, two- and three-dimensional transformations, parametric functions, window-viewpoint mapping, and parallel and perspective projections.

Spreadsheets are excellent tools for data visualization because they allow data to be easily input, dynamically updated, and readily viewed. Formulas can be created that include any combination of the spreadsheet's built-in functions, including trigonometric functions. These formulas are easily examined, tested, and changed. Their output values are automatically updated when the cells upon which they are constructed have their data values changed. Spreadsheets can contain dynamically linked charts, which are automatically redrawn when the data upon which they are based is changed. In addition, changing the position of data points in charts by using the mouse causes the data to be updated in the cells upon which they are based. Finally, spreadsheets, such as Microsoft Excel, are ubiquitous, included with the purchase of most computers. Therefore, spreadsheets can be thought of as a simple interactive graphic environment where data is input, equations are entered, and charts are built and interactively explored.

Using Spreadsheets
Spreadsheets are used to augment lecture demonstrations and as homework assignments. Incorporated into lectures, they are linked to derivations of formulas for two- and three-dimensional geometric transformations, window-viewpoint mappings, and projections. For lectures on parametric curves, spreadsheets are integrated into class presentations and discussions. For example, in a lecture about composite two-dimensional transformations, such as rotation about an arbitrary point in two-dimensional space, the students are led through the concepts of deriving a composite transformation with the sequence of steps: translation to origin, rotation about origin, and translation back. Matrices are

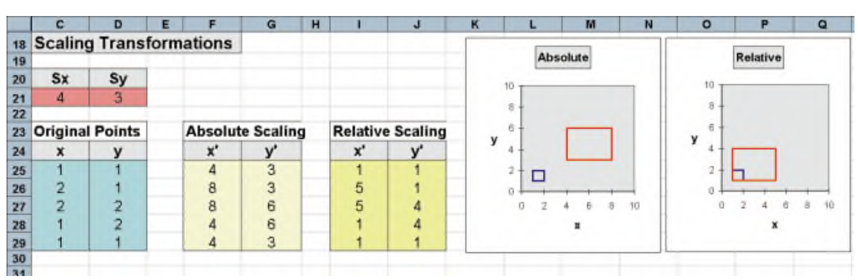

Figure 1 Excel spreadsheet demonstrating two-dimensional scaling transformations.

written on the blackboard and multiplied to create the final composite transformation matrix. Finally, the programmable equations are extracted from the matrix transformations. At this juncture, the instructor builds a spreadsheet using sample data, such as the unit square. Different angles are fed into the formulas and the output of each is discussed. At the end of the class session, students are handed screen dumps of the spreadsheet containing sample input data, output data, and the corresponding charts. They are expected to recreate the spreadsheet at home. The spreadsheet formulas that they have created become starting points for their programming exercises. The numerical output data become test cases.

Sample Applications

Four sample applications of spreadsheets are presented below: scaling transformations, rotation about an arbitrary point in the plane, parametric curves, and perspective projections.

Scaling Transformation

Figure 1 displays an Excel spreadsheet that has been constructed to demonstrate the effects of absolute and relative scaling transformations on geometric data. Absolute scaling not only enlarges the object but effects a translation, while a relative scaling transformation enlarges the object in place. This latter composite transformation results from a derivation where an object is translated to the origin, scaled, and translated back. The spreadsheet is laid out so the columns containing the original x,y coordinates are parallel with the x',y' counterparts that result from application of the transformation formulas for absolute and relative scaling. The highlighted cell (H25) has its formula displayed at the top of the figure, revealing its dependency on the x coordinate scale factor in cell C21, and the original x coordinate data in cell B25.

Excel's charting wizard is invoked to create two charts in the scatterplot chart format. Each scatterplot is constructed by Excel connecting each pair of data points in the sequence with a line. The coordinates of four points defining a unit square are placed in cells B25:C29 in clockwise order. The first point is replicated at the end of the list. This addition ensures that the polygon will be drawn with four sides. Each chart contains two polygons. The blue polygon is the original data. The red polygon is the transformed data. These charts are dynamic entities. They are automatically adjusted when there is a change in the contents of the cells from which they are built. If either scale factor is adjusted, formulas recompute the scaled coordinate values, and the chart is immediately redrawn.

The standout point about this spreadsheet is that it is an interactive textbook page. Instead of viewing static textbook images and relying on associated descriptive text, the student can alter input data to see the effects of both transformations side-by-side.

Rotation About an Arbitrary Point in the Plane

Figure 2 displays a spreadsheet for a 2D rotation of a unit square about a point x_c, y_c. As with the scaling example, an x, y scatterplot has been prepared to visualize the effect of changing the angle of rotation. The highlighted cell contains the transformed x coordinates of the first data point. Its equation is revealed at the top of the figure. Note calls to built-in sine, cosine, and pi functions. Note as well the conversion of the input angle from degrees to radians, a process mirrored in a student's implementation of the transformation equation in a programming language.

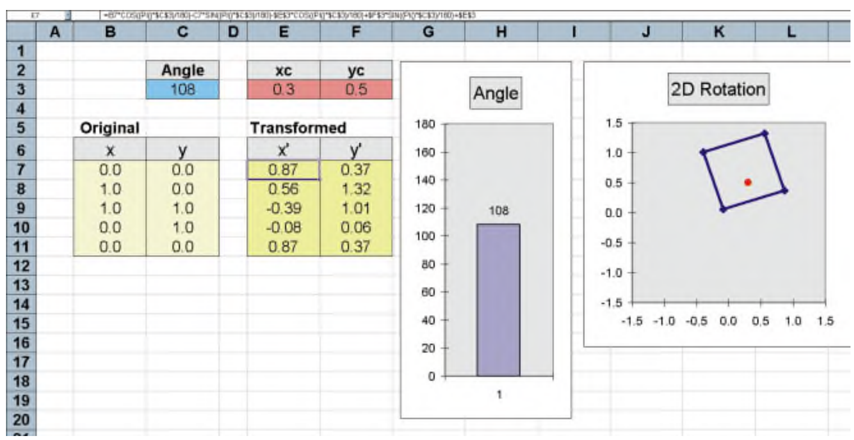

Figure 2 Excel spreadsheet demonstrating two-dimensional rotation about an arbitrary point.

It is possible to change a cell's contents by direct manipulation of its corresponding plotted point. The second chart found within this spreadsheet is a bar chart of the rotation angle. For example, by double clicking the mouse on the bar chart of the angle, and dragging its top boundary, the bar may be slid up and down. Angle data in cell C4 is simultaneously updated, which, in turn, updates the data points and their graphical representation.

Parametric Curves

One example of the strength of this method is in the teaching of parametric equations such as Bezier curves. The screen dump of a Microsoft Excel session is shown in Figure 3. The chart contains four control points connected by three straight lines and the smooth curve the control points generate. The control points reside in the cells above the plot. Computed x, y values for the smooth curve are in the cells to the plot's left. The spreadsheet was created in a matter of minutes with the definition of cells for control points and inputting of two formulas. The plot was created by highlighting the range of cells below the column headings P(x), P(y) and executing Excel's chart wizard.

The most interesting part of this spreadsheet is that new shapes of the Bezier Curve may be created by clicking on the control points in the plot and dragging them to any location. This process updates the control point data in the cells above the plot and in turn, updates the P(x), P(y) values.

Perspective Projection

An Excel screen dump from a lecture demonstration on perspective projection in presented in Figure 4. The original database defining a wireframe representation of a cube resides in cells C9:F16. The projected points are in cells H9:J16. Center-of-projection coordinates are found in cells H6:J6. Below the projected vertices are cells containing these points reassembled into front and rear polygons, followed by four line segments that connect them to complete the cube.

The Excel scatterplot for the perspective projection appears to the right of the wireframe data. It contains x_c, y_c, the x, y location of the observer, marked on the chart by a triangle, and the wireframe model of the cube.To create the wireframe image, it was necessary to group the datapoints, separating them by rows of empty cells. This was done because Excel connects each point in sequence. If left to its own devices, Excel's rendering would contain unwanted cross connections among points. Blank cells act as delimiters, telling Excel to stop connecting the dots. When a subsequent filled cell is encountered, Excel begins drawing anew. Hence, this data arrangement instructs Excel to draw two polygons and four line segments.

Below the original vertex data is a chart displaying a sideview rendering of the viewing environment containing data culled from all the spreadsheet's cells. From left to right, the chart contains: the viewer's position in the nega-

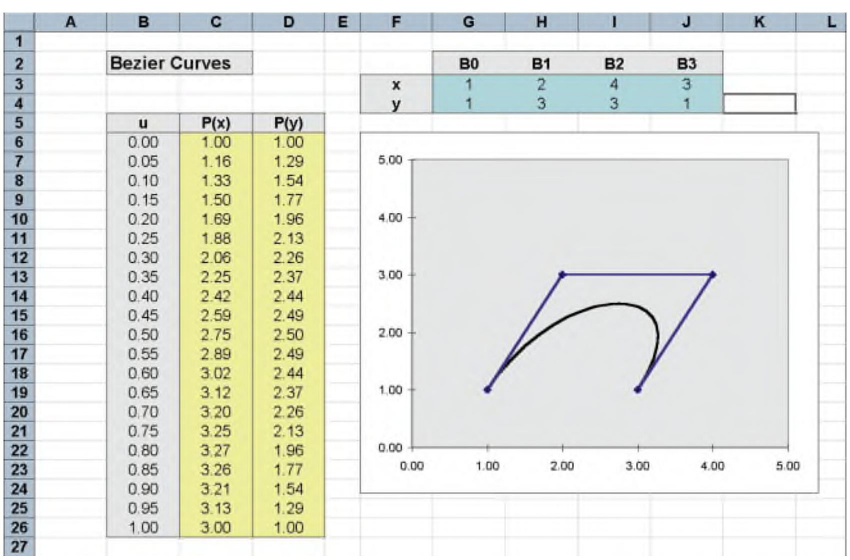

Bezier Curves

	B0	B1	B2	B3
x	1	2	4	3
y	1	3	3	1

u	P(x)	P(y)
0.00	1.00	1.00
0.05	1.16	1.29
0.10	1.33	1.54
0.15	1.50	1.77
0.20	1.69	1.96
0.25	1.88	2.13
0.30	2.06	2.26
0.35	2.25	2.37
0.40	2.42	2.44
0.45	2.59	2.49
0.50	2.75	2.50
0.55	2.89	2.49
0.60	3.02	2.44
0.65	3.12	2.37
0.70	3.20	2.26
0.75	3.25	2.13
0.80	3.27	1.96
0.85	3.26	1.77
0.90	3.21	1.54
0.95	3.13	1.29
1.00	3.00	1.00

Figure 3 Excel spreadsheet for Bezier curve.

tive z direction, viewplane projected points at z=0, and line segments depicting front and rear cube faces. This chart complements the perspective projection by showing the viewer's station point relative to the data and viewplane. It also shows the y values of each data point projected onto the viewplane. Notice that rear-face vertices are marked with filled diamonds and front vertices with filled circles. Most students have difficulty visualizing three-dimensional shapes. Their problems are compounded through projective transformation, which distorts these shapes in the projection plane. By using the spreadsheet to create multiple viewpoints, the student can better sense the dimensions of the visual space and the portions of this space occupied by the object, the viewer, and the projection plane. This spreadsheet works well as part of a lecture demonstration, where the instructor may move fluidly between object and projection changing parameters at will.

Conclusion

The strength of the spreadsheet approach is that the spreadsheet is an interactive visualization tool or scratchpad. Students can work out an understanding of their formulas and derivations, then build their programs. The use of spreadsheets can be easily taught as in-class demonstrations and integrated into the equation derivation process. The computer graphics topics amenable to the use of spreadsheets are: absolute and relative addressing, visualization of parametric equations including Bezier curves, window-viewport mapping, two- and three-dimensional transformations, and parallel and perspective projections.

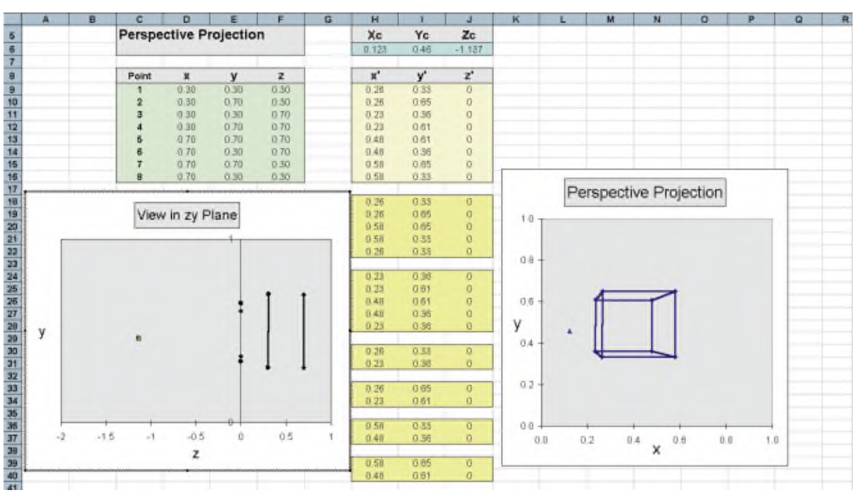

Figure 4 Excel spreadsheet for perspective projection.

Web-Based Teaching of Computer Graphics: Concepts and Realization of an Interactive Online Course

Reinhard Klein
Frank Hanisch
Wolfgang Straßer
Universität Tübingen

Abstract
Topics within computer graphics still cannot be adequately presented and explored with traditional teaching methodologies and tools. An integrative approach to combine lectures, examples, programming exercises, documentation, supported by a common sophisticated interface, is greatly needed. In this paper, we discuss and present the concept, realization, evaluation, and experiences of a computer graphics course that was developed at our lab to focus on this need. It is one of the first complete computer graphics course on the World Wide Web combining course text, programmed examples, and course exercises in a common hypertext framework.

Introduction
Visualization and interaction are the major topics of current computer graphics. Teaching these issues using only traditional teaching methodologies and tools, such as blackboards, slides, and even videos, cannot provide real-life examples. Only in real-life settings can the teaching react flexibly to the various questions that arise during discussions. As in physical research, problems and challenges in computer graphics can only be recognized through exploration in experimental situations. The experiments consist of programs.[19] Working with the programs not only helps to illustrate problems, but also increases student motivation. Furthermore, providing students with these programs enables them to repeat the experiments, build up their own settings, and deepen their understanding. This greatly improves learning.

In earlier times, such experiments were not included in traditional lessons. On the one hand, several different, expensive programs were necessary to cover the whole content of a computer graphics lecture, and, on the other hand, students didn't have appropriate hardware to run them at home. In addition, becoming familiar with all the different programs took too much time and could not be done in one or two semesters.

This also leads to serious problems with the practical exercises that accompany the lessons. A major goal of the exercises should be that students learn to implement the graphics algorithms discussed in the lecture. There were only two ways to achieve this goal: students built up their own graphics software systems from scratch or they used some of the existing program packages. Both were nearly impossible due to time constraints. The most common compromise was to present exercises that covered only a small amount of the teaching content.

A further major problem has been the number of different computing platforms and operating systems. It was nearly impossible to run the same program on all different platforms. Some of these ambiguities have been addressed by a number of colleagues working in the area of computer-supported education.[18, 2, 24, 12, 15, 17, 20] Many authors[23, 10] focus on the lack of a platform-independent graphics package that can be used as a basis for programming exercises. Fellner developed and used the object-oriented Minimal Rendering System (MRT) as a teaching and research platform for 3D image synthesis.[10] It is based on standard graphics packages like PHIGS, OpenGL, or XGL for graphics output.

Another approach is choosen by Wernert, [26 or 16] who describes a unified environment for presenting, developing, and analyzing graphics algorithms based upon IRIS Explorer, a modular visualization environment that provides a visual data flow language and allows users to link computational modules in order to create visualizations.

Lotufo and Jordan developed the well-known digital image processing system Khoros. A similar approach is described in.[16] This nice system applies high-level programming (network wiring) to more traditional coding of new modules. But there are also some drawbacks. First of all, IRIS Explorer is not platform-independent and is only available on high-end PCs and SGI workstations. Second, the application cannot be interlinked with a Web-based environment like our course. For students, even more interesting are the financial costs of a commercial system like IRIS Explorer, AVS or IBM Data Explorer.

In the meantime, the World Wide Web together with Java and Hypertext provide an appropriate framework to generate common interfaces for integration of all elements of interactive teaching courses, such as lectures, programs, exercises, and consolidating literature references. A certain number of computer graphics courses[4, 5, 14] already incorporate some of the ideas described above on the basis of HTML. The subsequent parts of this paper describe and evaluate the developed courses "Computer Graphik spielend lernen."[27] In the following section, we give a brief overview of the course, describe the Java-based programming environment and the compilation tools developed to integrate the different parts of the course, then consider how to generate reusable components using Java Beans. Finally, we conclude by reporting some experiences with the courses, describing our current work and giving a brief overview of our future activities.

The Course

Contents

The course contains the following topics: Computer graphics hardware, raster algorithms with aliasing and antialiasing, 3D transformations, visibility, color, local illumination schemes, modeling techniques, simple animation, texture mapping, global illumination techniques (ray-tracing, radiosity), basic image processing, and volume visualization. These topics are distributed through two courses. Both courses are based on, or related to, the books[8, 9] and corresponding scripts from the Fernuniversität (Correspondence University) in Hagen, Germany. The courses have been tailored to students of computer science and contain a variety of programming assignments.

In their current form, the courses are not well suited for engineering students: Such students prefer a course in which they can learn how to use commercially available packages instead of having to invest too much time in programming examples of basic computer graphics algorithms.

The programming exercises are done in groups of two or three students using the Java-based programming environment. The Java source code of the environment is available as a hyperlink and can be downloaded by the students. Comprehensive written instruction is available for each exercise in the form of HTML pages. In order to read it, students need to use a Web browser such as Netscape Navigator or Microsoft Internet Explorer. Since one of the aims of our course is to make the students familiar with graphics programming, it is important to provide them with sample programs to work with.

Structure

An HTML page provides a unified interface to our new Web-based computer graphics course. Starting at this page, one can follow links to HTML pages containing or referring to:

1. The instruction manual and editorial content of the course itself.
2. The script and list of available java-applets.
3. The programming exercises.
4. Links to external documentation and sources.

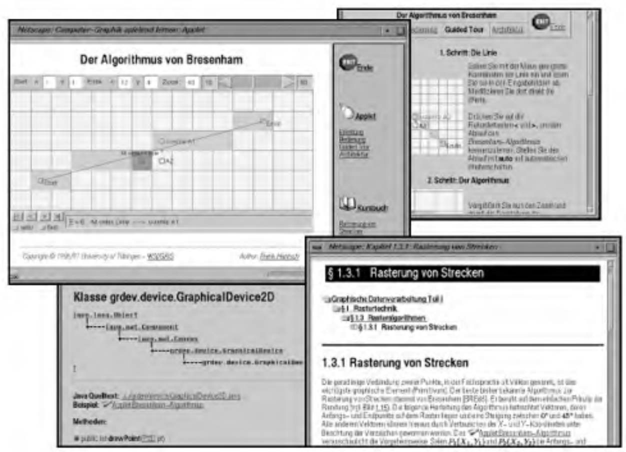

Figure 1 For each applet there is four-part documentation: An introduction, details on its features, a guided tour covering its essential topics, and general information on its programming architecture.

The instruction manual first gives a short introduction to the course and some hints for private studies, and provides a list of symbols used throughout the course. In addition, it gives a short overview of the design of the course, the programming architecture, and the file structures. Last but not least, it reports known bugs, such as the different behavior of certain applets on different browsers. This list of known bugs and comments can be extended by the user. External links provide tutorials for HTML and Java, public-domain software, other Web-based courses and further useful stuff.

The Script and the Java Applets

The entire course text is available in hypertext and presented in its own browser window. The contents and structure of the hypertext is the same as in the original course text. The hypertext contains not only cross references to figures, tables, literature, exercisesm and footnotes, but also links to corresponding Java applets and a number of videos and slides available via HTML that are shown during the lectures. If the user follows a link to an applet or a video, it is started in another browser window. The hypertext pages that contain the Java applets also provide links to the course text. In this way, the user working with an applet can immediately get the corresponding theoretical background. In addition, for each applet there is four-part documentation (cf. Fig. 1): An introduction, details on its features, a guided tour covering its essential topics, and general information on its programming architecture. Note that the course text, the applet, and its documentation are shown simultaneously in their own browser windows.

Web-Based Teaching of Computer Graphics: Concepts and Realization of an Interactive Online Course

Both the course text and the hypertext pages containing the Java applets contain links to the Java API. In this way, students have direct access to example implementations of the presented algorithms. The classes of the Java API contain links to applets that demonstrate their usage in real examples. For certain classes, like Camera, there are also links to the corresponding course text.

The Programming Exercises

For each programming exercise, a small example of the source code can be downloaded, along with a short introduction to the topic. The source code of the exercises is also accessible through the Web, thus allowing for easy switching among the theoretical background information, the description of the exercises, and the source code of the latter. The aim is to support the students during completion of their programming tasks. Via hyperlinks to the course text in HTML format, including the Java applets, students can review the theoretical background for every single exercise.

Implementation

Generation of the Hypertext

The ingredients for the hypertext were the original course texts as LaTeX sources with images, the Java applets, and the exercises. In the first step, we manually generated an applet resource file containing all information necessary to automatically generate a hypertext environment for an applet. This includes the four-part documentation of the applet and keywords used to generate links into the course text. The actual pages containing HTML tags for the applet, including the header, the links to documentation, the applet, and the hyperlinks to the source text, are then automatically generated by a Perl script. This guarantees an easy-to-modify and unified outlook of all applets and saves an immense amount of time in page design.

After the individual applet pages are generated, the applet resource files are used by another Perl script to automatically generate a page that contains an index of all applets. In addition to links to the applets, this page contains the title, introduction, and a motivating image for each applet.

In an independent step, the original LaTeX source files were modified. First, links to the different applets were placed into the course text. Then anchor points named by the header itself were automatically inserted. These anchor points are used later as hyperlink anchor points and are referred by the java applets. The names of the anchor points are copied into the corresponding applet resource files. After these two steps, we used the program Latex2html[6] to convert the course text into HTML format, which still contains the unprocessed anchor tags. These tags are processed in a further step by an additional Perl script, which generates the special course text structure developed for our course. This structure allows for folding and unfolding the hierarchy of the chapters. Each text page on the lowest level of the hierarchy automatically links to the previous and next text page as well as to all predecessors in the hierarchy.

The Java Applets

We put special emphasis on a common, easy-to-use interface for the applets. This is especially important, as the whole course includes a variety of different topics and the content of the functionality of the applets differs greatly. Some of the techniques we use include: using the same colors for the same context, using the same labels and control elements with identical meanings, and the same mouse control. A clear structure for the visible information is very important for design of the applets for teaching purposes. At first glance, students must be able to recognize the topic of the applet and the key elements of the teaching content and the connections between them. Therefore, the visual part of the applet containing this information must attract their attention more than special control elements that steer certain parameters. Otherwise, the applet's forest of equivalent choices would overwhelm the students.

The Graphics Engine

The graphics engine plays the central role in implementation of all the applets. Previously, Java has provided the user only with a 2D graphics interface. For real 3D graphics programming, we implemented our own 3D graphics engine. It supports all 3D functionality used throughout the applets, such as cameras, shading techniques, 2D and 3D texture, and hidden-surface elimination. It is based on our own 2D engine, which extends the standard Java graphics engine. This was necessary to supply abstract data structures, like 2D grids, and abstract event handling, like special mouse interpretations, as well as totally new components such as highlight modes.

The Scene Graph

The scene graph was developed as a programming basis for two reasons. First. it should provide a rich set of features for creating interesting 3D worlds and provide a high-level, object-oriented, programming paradigm that enables us to rapidly deploy sophisticated applications and applets. Second, it should reflect the state of the art of current graphics APIs[21, 1, 25, 7] in order to make the students familiar with the newest concepts of graphics technology. The scene graph is an acyclic directed graph and contains a complete description of the entire scene, or virtual universe. This includes the geometric data, the attribute information, and the viewing information required to render the scene from a particular point of view. The scene graph allows the programmer to think about geometric objects rather than about triangles, about the scene and its composition rather than about how to write the rendering code for efficiently displaying the scene. Applications construct individual graphic elements as separate objects and connect them together in the graph. If the Java3D[16] extension to Java is available, the basic part of our scene graph may be replaced by the corresponding components of Java3D, leaving most of our own parts intact.

Utilities

The basic mathematical and geometric material needed in most of the applets, like points, vectors, matrices, lines, circles, triangles, etc. is implemented as their own classes. This allows rapid implementation of even complicated mathematical and geometrical algorithms. Using these basic classes, the developer can concentrate on the essential parts of the algorithms rather than on basic mathematics. Furthermore, the resulting code for our applets becomes much easier to read and understand. This is especially necessary when motivating students to take a closer look onto the real implementation of an algorithm and try their own implementations.

The basic graphic components, like special interfaces for graphical in and output of the basic mathematical classes, the elements of scene graph, are realized in a similar way. A special component is a grouping panel to organize the contents of our course applets. Such components are used throughout the applets. Combining the standard Java graphics components with our own versions allows for rapid design of interfaces. Note that for all classes there exists hypertext-based documentation that can be invoked while browsing the source code and easily searched.

Reusable Components - Java Beans

Generating New Applets

The strict object-oriented design of our Java applets simplifies the programming of new applets and classes. Many of the existing classes can be reused and extended. Nevertheless, generation of new applets still requires low-level programming. Students or teachers not familiar with the existing classes have to wade through the jungle of hierarchical APIs instead of concentrating on the structure and concepts of the algorithms. As already shown by a number of authors, visual programming greatly aids in developing and debugging code.[26, 16]

The Programming Exercises

Since the low-level programming is too time-consuming without visual programming, our students only complete already existing applets. The drawback of this kind of programming exercises is that most of the students are not able to understand the overall structure of the program. Although they successfully implement the missing parts. there remains a bad feeling. Again, these problems can be solved using a visual application builder in combination with low-level programming and review of existing code.

The Beans

As stated in its specification,[11] a Java Bean is a reusable software component that can be manipulated visually in a builder tool. The components can be objects with or without graphical interfaces. In our course, typical beans are:

- Interface components that extend the standard Java AWT.
- Mathematical and geometrical utility components like vectors, matrices, triangles, spheres etc.
- 2D and 3D canvases that also extend the standard Java components.
- 2D and 3D scene graph for high-level graphics primitives and scene descriptions.
- Already-programmed “high-level” beans like image filters, function parsers, editable curves, etc.

Due to the strict object-oriented design of the Java classes used in the first version of the course, their conversion into Java beans was straightforward.

Examples

As in a visual data-flow language, beans can be visually composed into new customized applets. Figure 2 shows generation of a simple applet that demonstrates conversion of a gray-value image to a black-white image. The user can define the threshold value. To generate the applet, the image loader, two image beans for input and output image, the value panel, and the black-white filter are loaded into the application builder. In this example, the data flow is defined by so-called property-changed events. The property-changed event of the original picture (invoked by loading it) is propagated to the black-white filter that expects a reference to the input image and resolves a property-changed event with a reference to the output image. Last but not least, the property-changed event of the value panel changes propagates the threshold value to the black-white filter and invokes it. After the applet is designed in the application builder, it can be compiled into a new applet.

Figure 3 shows construction of a more complex applet demonstrating Bézier

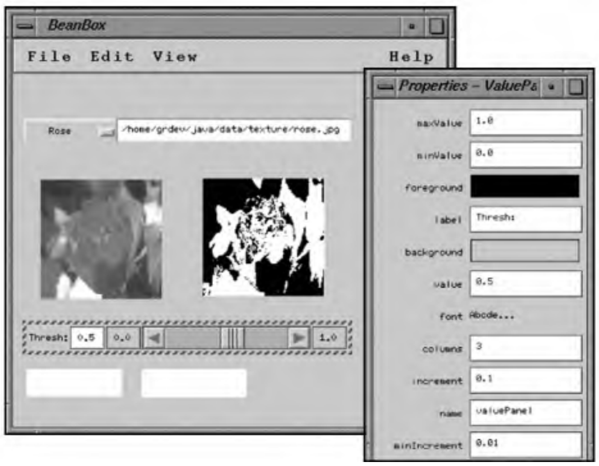

Figure 2 Building a simple applet demonstrating the black-white image filter in the Bean Box. Unfortunately, the current version of the Bean Box does not display the already-established links between the components.

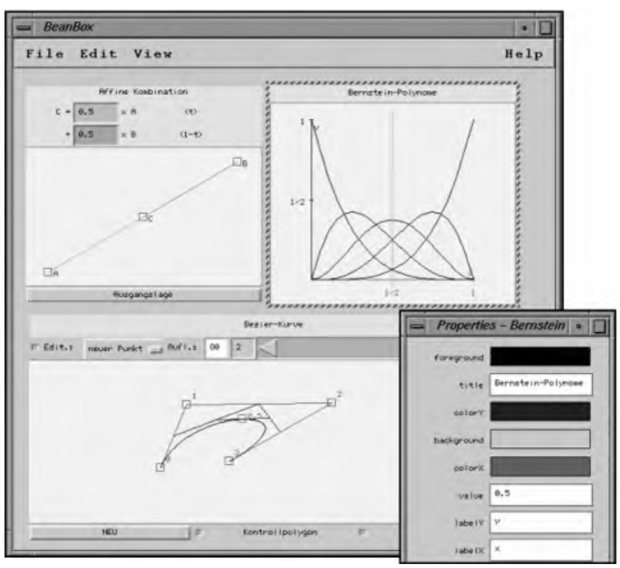

Figure 3 Composing three beans into a new applet demonstrating Bézier Curves together with affine combinations and Bernstein polynomials.

curves. Its main building blocks are three simpler beans, one realizing affine combinations between two points in 2D, one displaying the Bernstein polynomials, and one containing the Bézier curve itself. Each bean already implements all basic interactions to manipulate its contents. Several interface components link the beans.

In this way, the functionality of data-flow languages is available to build up new Java applications. Using this paradigm, teachers and students can easily develop their own new beans and integrate them into the course using existing components. Nevertheless, there are still cases, like the implementation of new filters, in which new beans must be implemented in the traditional way.

Discussion of Design Problems

One of the design problems of data-flow languages is how to manage the complexity of program networks. In our approach, this problem can (in most cases) be easily solved by grouping several beans into new ones and customizing their appearance within the builder tool. We made use of these techniques throughout the whole course. The problem that is inherent to most other systems based on the data-flow concept – propagating data through the network – does not occur, since different modules may share the same data in memory.

Experiences

Evaluation of the programming environment is based on two consecutive semesters of the computer graphics course offered in our department in 1995-96 and 1996-97. The complete first course consists of 28 hours of lectures and 26 hours of practical exercises; the second course consists of 24 hours of lectures and 20 hours of practical exercises. Forty-one students participated in the first course and 31 in the second. At the end of each course, the students were asked to fill out a questionnaire on their experiences with, and evaluation of, the complete course, the lectures, and the programming environment. The questions referred to issues like quality, effort with respect to results, learning success, and user content. The following trends could be traced:

- The overall content of the whole course was evaluated much higher than in the years before.

- The students themselves reported remarkable learning success.
- The complaints regarding the programming effort with respect to the outcome decreased.

We noticed that several students started to add non-requested features and functionalities to the system after they became familiar with it. These extensions referred to the user interface as well as the scene graph. Some students even implemented completely new applets in other topics in computer graphics that are not directly related to the central part of the lecture. This had never been the case before, and it strengthened our assumption that our concept would support and stimulate exploration and further development of the programming environment by the students.

Conclusions

Using our course, teachers and students can theoretically and practically prepare, catch up on, and deepen the course lectures. There is no need to provide for and install a variety of software packages to run the practical examples at home. Moreover, interested students are able to extend the examples. Due to the positive resonance in the course, we decided to refine and extend it to serve a larger community of people interested in computer graphics programming at universities, industry, and even high schools, and we are currently working on a translation of the whole course into English.

In addition to further development of our Java Beans toolkit, our future research in the area of graphics and visualization education has to consider appropriate user modeling, different levels of programming examples and even lessons, and, especially, different variants of course-related hypermedia user support.

Acknowledgments

We would like to thank A. Schilling for useful disscussions and G. Rössner and R. Schwering for their immense implementation efforts in realizing the system.

References

1 G. Bell, A. Parisi, and M. Pesce. The Virtual Reality Modeling Language: Version 1.0 Specification. URL: http://vrml.wired.com/vrml.tech/vrml10-3.html, 1995.

2 K. Brodlie, T. Hewitt, S. Larkin, P. Willis, and J. Gallop. Graphics and visualization - techniques and tools: A course for postgraduates of all disciplines. In 3., pages 263-268, 1994.

3 K. Brodlie and G. S. Owen, editors. Computer & Graphics, volume 18(3). Pergamon, May/June 1994. Special Issue on Computer Graphics in Education.

4 Course CPSC 453: Computer Graphics I.URL: http://www.cpsc.ucalgary.ca/localinterest/class info/453/, 1994. University of Calgary, Computer Scienc Dept.

5 Course Computer Science 417: Computer Graphics.URL: http://www.tc.cornell.edu/Visualization/Education/cs417/, 1996. Cornell University, Theory Center.

6 Nikos Drakos. The LaTeX to HTML translator. Internal report, Computer Based Learning Unit, University of Leeds, January 94.

7 George Eckel. Cosmo 3D Programmer's Guide. Silicon Graphics, Inc., 1997.

8 J. Encarnação, W. Straser, and R. Klein. Graphische Datenverarbeitung I. Oldenbourg, 4th edition, 1995.

9 J. Encarnação, W. Straser, and R. Klein. Graphische Datenverarbeitung II. Oldenbourg, 4th edition, 1997.

10 D. Fellner. MRT: A Teaching and Research Platform for 3D Image Synthesis. In 13., 1994.

11 Graham Hamilton. The javabeans api specification. Sun Microsystems, July 1997.

12 F. W. Jansen and P. R. van Nieuwenhuizen. Computer Graphics Education at Delft University of Technology. In 13., 1994.

13 L. Kjelldahl and J. C. Teixeira, editors. Eurographics Workshop on Graphics and Visualization Education (GVE), Oslo, Norway, 10-11 September. Eurographics, 1994.

14 Reinhard Klein and L. Miguel Encarnação. An interactive computer graphics theory and programming course for distance education on the Web. In 8th Int. PEG Conf.'97, Sozopol, Bulgaria, May/June 1997.

15 B. R. Land. Teaching computer graphics and scientific visualization using the dataflow, block diagram language Data Explorer. In S. D. Franklin, A. R. Stubberud, and L. P. Wiedeman, editors, University education uses of visualization in scientific computing: proceedings of the IFIP WG 3.2 Working Conference on Visualization in Scientific Computing, Uses in University Education, Irvine, CA, USA, IFIP Transactions. A, Computer Science and Technology, pages 33-36, pub-NH:adr, July 1994. pub-NH.

16 Roberto Lotufo and Ramiro Jordan. Digital image processing with khoros 2.0. This is an interactive WWW course using the Khoros system, December 1994.

17 Avi C. Naiman. Interactive teaching modules for computer graphics. j-COMP-GRAPHICS, 30(3):33-35, August 1996.

18 G. S. Owen. HyperGraph - A Hypermedia System for Computer Graphics Education. In S. Cunningham and R. Hubbold, editors, Interactive Learning through Visualization, pages 65-77. Springer-Verlag, 1992.

19 G. S. Owen. Teaching Computer Graphics as an Experimental Science. In 13., 1994.

20 Amnon Shabo, Mark Guzdial, and John Stasko. Addressing student problems in learning computer graphics. j-COMP-GRAPHICS, 30(3):38-40, August 1996.

21 P.S. Strauss and R. Carey. An object-oriented 3d graphic toolkit. In Computer Graphics Siggraph 92, pages 341-349. ACM Siggraph, July 1992.

22 Sun. The JAVA 3D API. Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303 USA., 1997.

23 J. C. Teixeira. Environments for Teaching Computer Graphics: An Experience. In 13., 1994.

24 J. C. Teixeira and J. S. Madeira. A Computer Graphics Curriculum at the University of Coimbra. In 3., pages 309-314, 1994.

25 N. Thompson. 3D Graphics Programming for Windows 95. Microsoft Press, 1996.

26 E. Wernert. A unified environment for presenting, developing and analyzing graphics algorithms. Computer Graphics, 31(3):26-28, August 1997.

27 Course Computer-Graphik spielend lernen. URL: http://www.gris.unituebingen.de/gris/grdev/java/index.html, 1996/97. University of Tübingen, Interactive Graphics Systems Lab (WSI/GRIS).

Distributed Development and Teaching of Algorithmic Concepts

Seth Teller, Nathan D.T. Boyd, Brandon W. Porter, and Nicholas J. Tornow
MIT Computer Graphics Group
graphics.lcs.mit.edu

Abstract
We describe Fuse-N, a system for distributed, Web-based teaching of algorithmic concepts through experimentation, implementation, and automated test and verification. Fuse-N is accessible to, and usable by, anyone with a Java-enabled Web browser. The system was designed to:

1. Minimize the overhead required for students to engage in the essence of the learning experience.
2. Allow students to experiment with, generate, and evaluate algorithms and their implementations.
3. Facilitate greater, more effective interaction among students and between students and teaching staff.

The system represents algorithmic concepts as dynamic "modules" in Java. Students map inputs to outputs either via system-supplied "reference" implementations; manually, with corrective feedback; or by writing one or more alternative implementations for the concept. The output of the students' implementation is programmatically and visually compared to that of the reference implementation. Modules can be interconnected in a dynamic dataflow architecture of arbitrary complexity. Teaching staff can monitor student progress online, answer questions, execute student implementations, and perform a variety of administrative tasks.

A prototype system supported two modules. We have since extended Fuse-N to include a classical polygon rasterization pipeline, a module authoring "wizard," an editor, and several other components. We describe Fuse-N from student, staff, and developer perspectives. Next we describe its architecture and the technical issues inherent in its construction and extension. Finally, we report on some early experiences with the system.

Motivation
Computer science courses with large enrollments and significant programming components stretch the abilities of teaching staff to provide effective infrastructure, and individualized attention, to their students. For example, students need development tools and standardized (or at least consistent) environments in which to generate, test, and submit their implementations. Providing and maintaining these can be a formidable engineering task for the staff. Students submit complex programs requiring evaluation by the staff, often reducing the time available for individual attention to students. Also, despite nearly ubiquitous networks, email, etc., it is difficult, with standard mechanisms, to ensure effective, timely communication and collaboration among students and between students and staff.

The Fuse-N system addresses each of these considerations. First, it provides an intuitive, consistent visual interface through which students can experiment with, generate, and test algorithm implementations, ultimately submitting them for grading. This interface is backed by a robust client-server architecture for maintaining persistent versions of assignments and student implementations. Second, Fuse-N provides an interface through which the teaching staff can assess student work in progress and submitted work, using a location-independent mechanism for executing students' implementations. A development kit for generation of new modules is also provided. Third, the system provides mechanisms for synchronous and asynchronous communication among students and between students and staff, including annotations of course material, chat boards, shared whiteboards, and staff grading and feedback for specific student work.

Previous Work
The Web-based educational tool WebCT [Gol+96] provides sophisticated authoring tools to embed general course content, but no specific support for teaching algorithmic concepts is included. Instructional software development environments are less common, though the Jscheme environment [Hic97] allows students to implement assignments in a subset of the Scheme language from any Java-enabled Web browser. Algorithm visualization environments and toolkits, such as the Zeus system [Bro91] and Brown University's Interactive Illustrations [Dol97,Try97] provide sophisticated frameworks for examination of running algorithms. Finally, dataflow systems such as SGI's ConMan [Hae88], IBM's Data Explorer [IBM98], and AVS [Ups+89] allow direct manipulation and aggregation of functional modules. (An in-depth discussion of these and other related systems can be found elsewhere [Por98].) To our knowledge, however, Fuse-N is unique in its combination and synthesis of many such elements into a system for distributed collaborative pedagogy.

System Users' Perspective
One effective way to view the Fuse-N system is from a series of user perspectives. The system is first described from the point of view of a student. The teaching staff's point of view is described next, including mechanisms for disseminating course material, and collecting and evaluating student work.

Finally, the developer's perspective details the steps required to author a new module for the system. Our discussion largely avoids implementation details; these are deferred to Section 4 (System Designer's Perspective).

Student Perspective

Unorganized, the sheer volume of available course material can overwhelm a student. Thus the educator has two principal tasks: first, to select a collection of material to be taught, and second, to organize this material into a coherent progression of concepts. Fuse-N facilitates exactly this sort of progression in its exposure of concepts at several levels of abstraction (or, equivalently, at several levels of detail). For example, computer science and engineering courses typically emphasize the functional composition or interrelation of a collection of smaller pieces. Fuse-N represents such building blocks as algorithmic components to be manipulated by the student. These components are visible as modules within the Concept Graph, a gridded workspace containing a persistent collection of interconnectable modules selected by the course staff (Figure 1). Each module represents an algorithmic concept: a well-defined algorithmic mapping of input instances to outputs. Each module has a visual representation, consisting of an icon, a text label, multiple selectable input and output ports, and an interface with which the user can select among the interaction modes defined below.

In line with the usual notion of abstraction, students can think of modules and their connections either as closed functional units, or as open, modifiable components. For example, students can connect one module's output port to another module's input port (provided their types match), thereby assembling an aggregate functional unit. Students can open a module for editing, supplying their own code to implement the module's function. Successive student implementatinos can be compiled and reintegrated into the running system on-the-fly (Section 4.1).

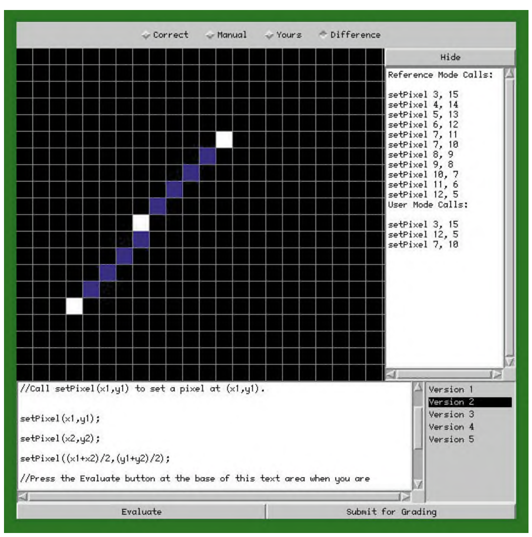

Figure 1 The Fuse-N work area

Students who actively engage in implementing an algorithm understand and retain concepts better than those who simply, and passively, view visualizations [Law+94]. In Fuse-N, students can assess their understanding with the help of four interaction modes: Reference, Manual, Implement, and Difference. In all cases, algorithm inputs are either provided by the system or specified interactively by the student or staff. Regardless of mode, a diagnostic area reports a textual description of each input instance and each output action.

Reference Mode allows the student to view the operation of a correct "reference" implementation (cf. Figure 1) provided by the module developer (Section 3.3). The source code for the reference implementation is generally hidden from the student, though in some cases it may be exposed (for example when the student is challenged to implement some alternative, or asymptotically superior, strategy). Reference mode is similar to that provided by traditional algorithm animation and visualization systems [Bro+84]; it helps a student understand algorithm behavior for a variety of input instances, including boundary cases in which the correct behavior might not be otherwise evident.

Distributed Development and Teaching of Algorithmic Concepts

Manual mode challenges the student to effect the algorithm's operation solely through user-interface actions (Figure 2). Students need not do so by simulating an algorithm; their task is only to produce the correct output for the currently specified input. The system provides an immediate visual response to each student action. For instance, Manual mode for Bresenham's algorithm challenges the student to select, for a given line segment specified by its endpoints, those pixels which would be drawn by Bresenham's segment rasterization algorithm [Fol+82]. Correctly placed pixels are drawn normally, whereas incorrect pixels appear red. Manual mode inputs are typically crafted to allow students to generate the correct outputs in a few minutes.

Implement mode challenges the student to supply Java code that implements the algorithm for the current module (Figure 3). In this mode, the student is effectively presented with the interface to a Java class, whose methods are invoked by the system during interaction. Continuing with our example, Implement mode provides a stubbed method Bresenham(x1,y1,x2,y2), which is invoked whenever the student selects or specifies a point pair, or a suitable input event arrives from some upstream module. The interface also provides one or more functioning "base calls," which map to the module's abstract output ports. Thus the base call SetPixel(x,y) effects the illumination of the pixel at framebuffer position (x,y), causes the printing of a diagnostic message, generates a typed event for forwarding to downstream modules, or performs any combination of these, according to the current interaction context.

Difference mode displays a visual representation of the difference between the output of the student's implementation and that of the reference implementation (Figure 4). For example, difference mode for the Bresenham module colors pixels in one of four ways. Pixels drawn by both the reference and student implementation are shown in white; those omitted by the student algorithm are shown in blue; those drawn in the wrong position are shown in red; and those output more than once are shown in green. Difference mode is a particularly useful and unique feature of Fuse-N; it allows the student (and staff) to determine quickly the algorithm's correctness for a variety of prefabricated or interactively specified input instances. (Crafting efficient and pedagogically useful difference engines is a challenging part of developing a new module; we address this issue later.)

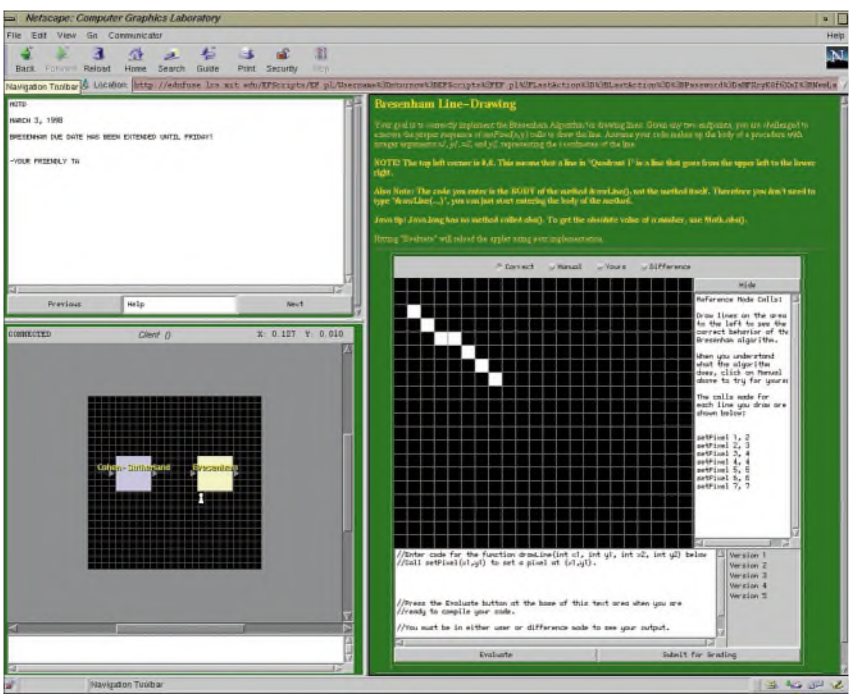

Figure 2 Manual mode

The student can invoke the Fuse-N editor for any module. This is a rudimentary editor that displays the class interface, staff comments and hints, and the student's code. It also allows the student to load, save, and compile implementations, and view compiler diagnostics. The editor will eventually incorporate debugging operations, student and staff annotations, and other administrative functions (see Current Efforts and Future Directions).

A suite of communication tools round out the student's view of Fuse-N. Students can read the message of the day, chat with other students or teaching staff who happen to be online, browse the chat history to look for questions that may already have been asked, or share a graphical whiteboard. We are extending these features to allow "sharing" of a client's workspace.

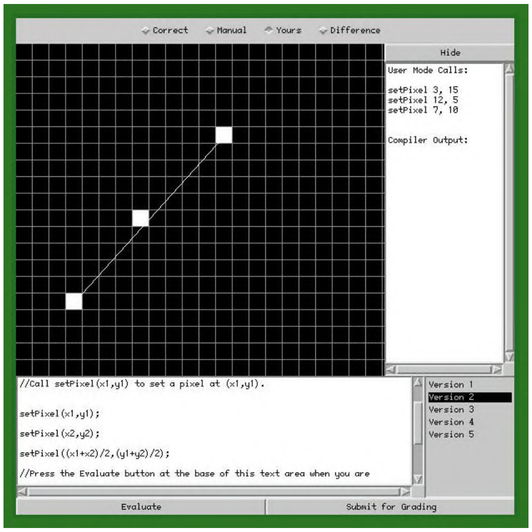

Figure 3 Implement mode

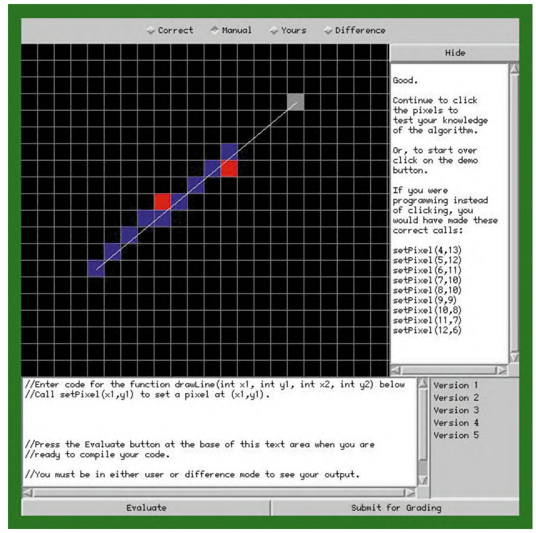

Figure 4 Difference mode

Thus students and staff will be able to interact with a Concept Graph and modules as if they were at the same console; for example, a teaching assistant could demonstrate a certain input case to the student by direct manipulation of the student's desktop. This and many other collaboration aspects of Fuse-N are discussed elsewhere [Boy97].

Teaching Staff's Perspective

Fuse-N is designed to maximize the teaching staff's availability and effectiveness through provision of infrastructure for creation, selection, distribution of course material, tracking and evaluation of student work, and general course administration such as maintenance of course lists, lecture notes, grades, daily messages, and other materials.

Generating, testing, organizing and distributing course software infrastructure to students can be a significant burden for teaching staff. Fuse-N provides a simple mechanism for authoring modules that represent algorithmic concepts, and functionally interconnecting these modules within a dataflow architecture with visible structure. Related materials are available through standard Web mechanisms like linking.

One traditional coursework model involves the staff's generation and assignment of work: assignment completion; a deadline at which time students turn in completed work; an evaluation period for staff to process this work; and a feedback period in which a grade and other information are returned to the student. These phases routinely span several weeks of activity.

Fuse-N provides several mechanisms that extend this batch model to an interactive setting. For example, the system places a visual avatar for each student near the visual icon representing the module on which the student is working (Figure 2). The teaching staff can broadcast timely messages to selected groups. The staff can also monitor an individual student's progress by locally executing the student's current implementation in Implement or Difference mode as described above. An assistance queue, interactive chat sessions, whiteboards, and this "over-the-shoulder" execution mode all increase the staff's ability to identify struggling students and help them in a timely fashion.

Several mechanisms are provided to ease the process of evaluating large amounts of student code. The staff can apply prefabricated test cases to student implementations. Common errors are then identified by interposing the student module between a generate/test module pair. A Web-based grading form (Figure 5) prompts the staff member for specific feedback about the student work and can incorporate input instances that cause the student implementation to exhibit specific behavior. These mechanisms log grading information and associated material to the server for persistent storage.

Finally, the system provides a uniform, location-independent interface to all of these mechanisms. Thus just as the students may complete coursework from any browser, so can the staff grade or otherwise provide feedback to any student from any browser (Figure 5). We have found this capability to be of significant value in actual course administration (Section 5.1).

Creating a new module is a relatively simple process. Using common design patterns, the basic functionality of every module can be coded automatically. Through a dialog box, the developer specifies the module name, the input type(s) that the module can receive, and the output type(s) that the module generates. A Java servlet, analogous to a development kit “wizard,” generates a Java source file representing the module and compiles it into a class file. The Fuse-N client then requests that the class be loaded into the current application, creating an editable instance of the new module.

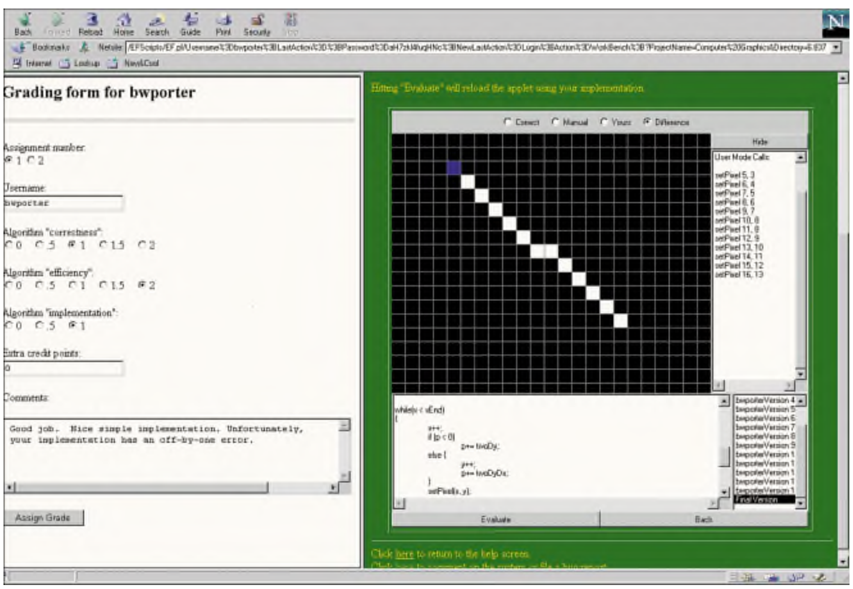

Figure 5 Grading a student from within a browser

Module Developer’s Perspective

This section describes the system from the point of view of the module developer, who maps abstract algorithmic concepts to actual modules in Fuse-N. Module developers may be teaching staff, students, or more broadly any third party who wishes to generate course content.

In order to implement a specific module, the module developer first provides code that maps both input and output data to visual representations on a graphics canvas associated with the module. Next, the developer codes the actions to be taken when new data arrive on any input port. The developer issues data on an output portsimply by invoking a provided method call. The developer must also design, specify, and provide the base call(s) with which the student is to generate output. Each base call must then be mapped to the appropriate output port. Finally, the developer must specify and implement the module’s behavior in Manual, Reference, and Difference modes. Once a module has been created, it is easily added to the Fuse-N server, becoming available to all clients of that server.

System Designer’s Perspective

This section describes Fuse-N from the point of view of its designers. Here we discuss the technical underpinnings of the system and some of the issues inherent in its construction and extension.

System Architecture

This section briefly outlines the technologies used in the system architecture. Specifically we discuss our choice of a client-server model, the advantages of location- and platform-independence, and Fuse-N’s dataflow architecture.

Fuse-N is implemented as a client-server architecture. The Fuse-N server is comprised of a traditional Web server and a group of Java servers and Java servlets. A Java registry server tracks users in the system; a Java messaging server handles text-based message delivery; and a Java servlet performs essential system functions such as compilation of user modules and generation of HTML code for display to the client. All user, course, and system information is stored persistently on the server. The server supports location-independent access to user data, restricting access with a simple password mechanism.

The client application is comprised of an HTML environment and a set of Java applets, both of which are served to the user's browser. Users receive the most recent version of the Fuse-N client at the start of each session. Because Fuse-N is coded entirely in Java, Fuse-N clients and servers can run on any platform equipped with a suitable Java Runtime Environment [Gos95].

Fuse-N uses a stateful dataflow model of sequential execution, enabling arbitrary aggregation of components. Software applications can be created by linking the output(s) of stateful components to the input(s) of other components. Thus an application can be modeled as a directed acyclic graph of components. Because the input(s) and output(s) of the components in a dataflow system are defined explicitly, each component can be treated as a black box abstraction of the operation it performs. The resulting dataflow aggregations are modular and can be rapidly assembled, modified, and extended. Dataflow systems naturally extend to parallel operation, as each component processes data independently of all other components.

Fuse-N's underlying architecture supports the ability to load modules, manage interconnections, and route information from any output to any input. Thus users create complex systems by linking modules in a run-time dataflow system, as in ConMan [Hae88]. Individual components have graphical representations that can be arranged and linked by the student. Links are typed and can be created, modified, or destroyed.

Modules run in separate threads. Each module queues incoming data for processing and outgoing data for delivery to downstream modules. Modules can be recompiled dynamically; successfully compiled modules are inserted into the running dataflow. Newly created modules inherit all of these capabilities through a subclassing mechanism. The student is left to focus entirely on the module algorithm, interconnections, and graphical display. (A comprehensive overview of Fuse-N's dataflow architecture can be found elsewhere [Bwp98].)

Early Experiences

Two specific experiences with Fuse-N have been illustrative. First, we set out to assess the system's effectiveness by deploying two simple modules as assignments for an undergraduate computer graphics class at MIT. Second, we implemented a "classical" polygon rasterization pipeline as a series of dataflow components within Fuse-N. After reviewing these experiences, we describe several current and future directions for Fuse-N.

A Trial Run

Fuse-N's early prototypes were designed primarily as proofs of concept. Fuse-N was first deployed to students in the fall of 1997, in the form of an introductory computer graphics course with 104 students. Each student was asked to implement two modules inside Fuse-N: Bresenham's segment-drawing algorithm and Cohen-Sutherland's segment clipping algorithm [Fol+82]. Upon submission of the assignment, each student was asked to fill out an anonymous survey. The results were positive; students felt the assignments were well organized and helped them learn the material.

As expected, human interaction was a crucial element of the process. The collaboration tools of Fuse-N were heavily used, allowing teaching assistants to pinpoint many of the errors students were making only minutes after students began coding. With the students' permission, teaching assistants effectively provided preemptive assistance. They did not have to wait for students to start asking questions, but could load students' code and examine its behavior. The staff could then clarify technical issues and make suggestions to the entire class, or only to those students working on a specific module. This immediate problem detection and correction ability allowed students and teachers to interact effectively outside of the classroom.

Another interesting aspect of the system was the staff's ability, given access to the development process as well as the final results, to distinguish among the learning and implementation styles of the students. Some students wrote careful pseudocode, then a small number of implementations. Others seemed to be coding by successive approximation, submitting many scores of slightly varying code revisions until a correct implementation was finally achieved. This rather unexpected result - that in removing essentially every barrier to recompilation we may have actually encouraged an inferior programming practice - has led us to reexamine the way in which students are asked to use the system. In particular, we are now investigating mechanisms by which students can be urged to achieve greater mastery of the material before attempting implementation. For example, students could be prompted to demonstrate understanding of the input/output mapping in Manual mode, and to submit coherent pseudocode to a staff member, before commencing implementation.

Implementing a Polygon Rasterization Pipeline

We extended, and successfully tested, the system by implementing a classical polygon rasterization pipeline from a core set of modules (Figure 6) including TraverseScene, CameraModel, BackFaceWorld, WorldToEyeSpace, EyeSpaceToNDC, ClipNDC, NDCToScreenSpace, RasterizeTriangle, DepthBuffer, and FrameBuffer. Each member of the development team implemented several modules. Instantiating each module required clicking the New Module button, entering the module's name, and specifying the module's typed inputs and outputs. A compilable stub for the module was then automatically generated and compiled by Fuse-N. Finally, we filled in the stubs by writing Java code to effect the functional mapping required of each module and wired together the appropriate module inputs and outputs through the Concept Graph interface. Implementing this system took three students about one week. This was less time than we had anticipated. (However, the students implemented only the Reference algorithm for each module, arguably the easiest portion; the Manual and Difference mode implementations are underway.) The system proved robust, supporting multiple simultaneous execution paths while maintaining responsiveness. The dataflow model was a natural fit to the graphics pipeline. Students who implemented parts of the pipeline reported that even though they previously understood the pipeline, they had a much better sense of it after implementing it in Fuse-N. The students also learned that adding internal, "self-checking" routines to modules was a particularly useful technique. Overall, the Fuse-N environmentmade the process of building and exercising this system efficient and educational.

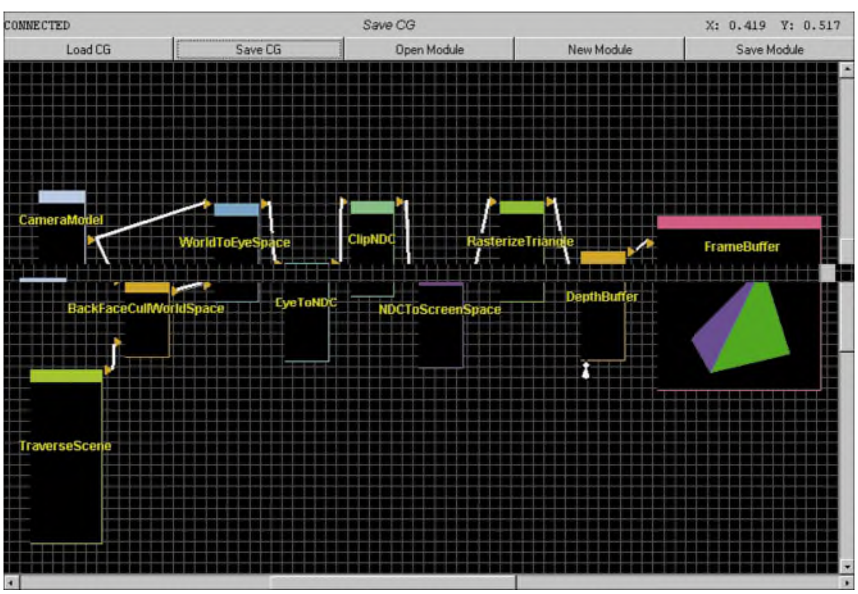

Figure 6 A Dataflow Polygon Rasterization Pipeline

Current Efforts and Future Directions

We continue to add capabilities to the system. Our active research and development efforts can be categorized as: improving the integrated environment for students and developers; crafting new capabilities for the teaching staff; and extending the dataflow model.

The current editing and debugging capabilities are adequate for simple modules but ill-suited for more complex implementation tasks. We are actively incorporating an improved code editor and debugger into the system. Mechanisms for more closely coupling applets, dynamic module state, and compiler and run-time diagnostics are being explored as well.

We are also expanding the capabilities of the teaching staff by improving the system's ability to track students. More generally, we are enhancing Fuse-N's general administrative features. For example, the staff expressed the desire to revise grading criteria, generate Web-based grading forms, and analyze and distribute grades from within the system. Finally, we have designed for transparent sharing of events across the network [Boy97]. This will allow teaching assistants to effectively reach over a student's shoulder and assume control of the mouse or keyboard, providing a kind of remote tutoring.

A related thrust of our work is extending modules to connect to other modules on remote machines, using a transparent network event layer [Boy97]. Providing basic dataflow constructs will help students and teaching staff assemble complex systems easily. Abstracting these connective elements as components in their own right is a notable innovation of our system. Finally, determining how best to integrate manual interaction with the dataflow system and accompanying visualizations, and how to simplify generation of difference engines for complex algorithms, are both hard open questions.

We foresee developing a number of additional modules as part of the Fuse-N project to augment the full undergraduate computer graphics curriculum currently under way. We are exploring collaborations with other algorithmic learning domains, such as those of MIT's class 6.001, The Structure and Interpretation of Computer Programs [Abe+85]. There, we plan to support a Scheme interpreter within a Fuse-N module, allowing dataflow interconnection of "Schemelets." We hope that as Fuse-N becomes more widely used, students, educators and researchers will contribute modules to the system, creating a distributed library of instructional algorithmic modules in the public domain.

Conclusion

Fuse-N improves upon traditional software development environments, while extending Web-based educational systems to support experimental problem solving. We designed it explicitly to support interactive pedagogy of algorithmic concepts. The system supports rapid engagement of students with the essence of the material. Fuse-N supports algorithmic instruction through Reference, Manual, Implement, and Difference modes for each module. At a higher level of abstraction, students can link modules in a dataflow environment. Preliminary results are encouraging, as gauged both by classroom experiences with an early prototype and by our nascent realization of a classical polygon rasterization pipeline within Fuse-N.

Acknowledgments

We gratefully acknowledge the past and present Fuse-N development team, including the authors: Bhuvana Kulkarni, Aaron Boyd, Randy Graebner, and Arjuna Wijeyekoon. The Fuse-N project has been supported by the MIT Class of 1951, the MIT Lab for Computer Science, the National Science Foundation (through Career Award IRI-9501937), Silicon Graphics, Inc., and Intel Corporation.

References/Bibliography

[Abe+85] Abelson, Harold and Sussman, Gerald and Sussman, Julie. Structure and Interpretation of Computer Programs. MIT Press/McGraw-Hill, 1985.

[Boy97] Boyd, Nathan. A Platform for Distributed Learning and Teaching of Algorithmic Concepts. MIT MEng Thesis, 1997.

[Bro+84] Brown, Marc and Sedgewick, Robert. A System for Algorithm Animation. SIGGRAPH Proceedings: Volume 18, Number 3. July 1984. Pages 177 - 186.

[Bro91] Brown, Marc. Zeus: A System for Algorithm Animation and Multi-View Editing. In IEEE Workshop on Visual Languages, pages 4-9, October 1991.

[Dol97] Dollins, Steven. Interactive Illustrations. http://www.cs.brown.edu/research/graphics/research/illus/. Brown University Department of Computer Science, 1997.

[Fol+82] Foley, James, and van Dam, Andries, Fundamentals of Interactive Computer Graphics. Addison-Wesley, 1982.

[Gol+96] Goldberg, Murray et al. WebCT - World Wide Web Course Tools. http://homebrew.cs.ubc.ca/webct/. University of British Columbia Department of Computer Science, 1996.

[Gos+95] Gosling, James and McGilton, Henry. The Java Language Environment: A White Paper. Sun Microsystems. October 1995.

[Hae88] Haeberli, Paul. ConMan: A Visual Programming Language for Interactive Graphics. SIGGRAPH Proceedings: July 1988. Pages 103-111.

[Hic97] Hickley, Tim. Jscheme: an Applet for Teaching Programming Concepts to Non-Majors. http://www.cs.brandeis.edu/~tim/Packages/Jscheme/Papers/jscheme.html Brandeis University, Michtom School of Computer Science, 1997.

[IBM98] International Business Machines, Inc. IBM Data Explorer (DX), http://www.almaden.ibm.com/dx/, 1998.

[Law+94] Lawrence, Andrea, and Stasko, John, and Kraemer, Eileen. Empirically Evaluating the Use of Animations to Teach Algorithms ftp://ftp.cc.gatech.edu/pub/gvu/tech-reports/94-07.ps.Z Technical Report

GIT-GVU-94-07. Georgia Institute of Technology College of Computer Science. 1994.

[McC95] McCanne, Steve, and Jacobson, Van. vic: A Flexible Framework for Packet Video, ACM Multimedia '95.

[Por98] Porter, Brandon. Educational Fusion: An Instructional, Web-based, Software Development Platform. MIT MEng Thesis, 1998.

[Try97] Trychin, Samuel. Interactive Illustration Design. Brown University Computer Graphics Lab, 1997.

[Ups+89] Craig Upson, Thomas A. Faulhaber, Jr., David Kamins, David Laidlaw, David Schlegel, Jeffrey Vroom, Robert Gurwitz, and Andries van Dam. The Application Visualization System: a Computational Environment for Scientific Visualization. IEEE Computer Graphics and Applications, Vol. 9, Num. 4. p.30-42. July 1989.

Enhanced Realities

Introduction

"Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!"

The Red Queen

Lewis Carroll, "Through the Looking Glass"

In his tales of Alice's adventures, Lewis Carroll uniquely captured a world of wonder as seen through the eyes of a child. Even as adults, we can relate to Alice and the nonsensical creatures and situations of Wonderland. Perhaps we more than half believe that there are truths to be found there. But too often the rest of our "grown up" perception sees the world as all too ordinary, which may point to nothing more than a suspension of wonder. It's all too easy to lose our ability to ask Alice-like questions, to wonder: "why?"

We hope that your walk through the contributed wonderland of fresh ideas that created Enhanced Realities will rekindle your sense of childlike wonder, the essence that is SIGGRAPH. Our program is devoted to those who continue to ask why, and to those in whom inspiration and wonder has been renewed in their asking.

Of the more than 50 proposals submitted to Enhanced Realities, we chose this year to accept just the top layer, the 17 most impressive and groundbreaking innovations, for presentation at SIGGRAPH 98. This work envisions our augmented future with clever multimodal interfaces that challenge our ideas about computing in the physical world and question this dubious concept called "reality." The goal of Enhanced Realities is to make us return to that childlike wonder of discovery, to inspire us with technological innovations that immerse us in a new, enhanced reality.

I am very grateful to all the wonderful world of potential contributors – the whole of the SIGGRAPH community and beyond – whose work points the way to our most creative emerging achievements on the horizon of technology. Many valuable lessons and insights arose in our Enhanced Realities jury meeting, including several outstanding inspirations for new projects based on the work our contributors submitted!

In the following pages, we share some "insider information" on how the selection process works, what a jury looks for when reviewing proposals, and how you can best present your vision and achievements to your busy but excited peers.

Janet McAndless
Chair, Enhanced Realities

Sony Pictures Imageworks
9050 W. Washington Boulevard
Culver City, California 90232 USA
Janet_McAndless@siggraph.org

Enhanced Realities Review Jury

Chair
Janet McAndless
Sony Pictures Imageworks

Jury
Bruce M. Blumberg
MIT Media Lab

Andrew Glassner
Microsoft Research

Jacqueline Ford Morie
Blue Sky | VIFX

Ken Perlin
New York University

Kathryn Saunders
Royal Ontario Museum

Enhanced Realities Committee
Bruce Dobrin
Janet McAndless
Preston J. Smith
Michael Trujillo
Alex Whang

Enhanced Realities is dedicated with heartfelt thanks to:

All those who took the time to submit an emerging technologies proposal, our contributors who undertook the time and expense required to show us their visions of the future, and our jury participants whose careful review efforts greatly enhanced our program.

Walt Bransford and Dino Schweitzer for unfailing encouragement and guidance.

Sony Pictures Imageworks for their uncommon support, including use of facilities and resources.

The Enhanced Realities Committee for their vision, dedication, and labor.

The SIGGRAPH 98 Conference Committee whose indispensable help and feedback formed our enhanced reality.

Jury Statement **Bruce Blumberg** MIT Media Lab

What I was looking for (and sometimes found)
As a jury member, I asked four big questions about each submission:

- Is it clear?
- Is it fun?
- Is it real?
- Is it important?

Is it clear?
It would be nice if the jury members spent as much time reading your proposal as you undoubtedly spent writing it. The sad fact of the matter, however, is that your submission is likely to be one of 100 that the jury put off reviewing until several days (hours) before the jury meeting. So it is absolutely essential for you to present your proposal as clearly and as succinctly as possible. Clearly articulate what the user experience will be, and the technical significance of the piece. If it seems similar to previous installations, then be sure to put it in context (be clear about what is different). Try to do all of this in a page or two, and in your accompanying video. Have a friend who doesn't know your work read it and tell you what you are proposing to do, and see if they get it. You are only hurting yourself if you make the jury struggle to understand your proposal. Struggling juries are grumpy juries!

Always ask yourself how you can make your proposal stand out from the 100 other proposals. This is often as hard as the underlying implementation, but it is well worth the effort.

Is it fun?
I am most attracted to submissions that seem like they would be fun to experience. Fun does not in anyway imply that the underlying work is not serious and important, but rather that it is presented in a way that is accessible and clear to the guest. Indeed, "fun" is a good indicator of a great implementation and/or a novel idea. Submissions should make it clear through the video and description why an installation will be fun to experience. A great implementation based on a "little idea" can often be a more compelling installation than a "big idea" done poorly. This is not to say you should focus on little ideas, but if you have a little idea (not Turing-award material) and you know how to turn it into a fun installation, go for it. Chances are the jury will go for it too.

Statements that lapse into "art criticism speak," or videos that seem as if they were done with MTV in mind come off as pretentious and work against you. The jury isn't exactly "hard-boiled," but it is composed of people who spend more time building things than talking about things. They want to hear how an installation will feel to the guest, so they would rather read a discussion of the underlying technology than learn that the installation is a "post-modern deconstruction of (whatever)." The word "chaos" is a good word to avoid. It has been known to cause rashes among more than a few jury members.

Is it real?
You should go out of your way to encourage the jury to believe that there is some chance that you can pull off what you say you are going to pull off. For example, if your installation is predicated on natural language understanding, the jury is going to be very interested in understanding how you intend to solve that problem. Videos, published papers, and the past experience of your team can all help to bolster the jury's confidence in your abilities. Don't be afraid to add a "this is how it works" section to your proposal. This is particularly true if you are pushing the state of the art.

Speaking of videos, do yourself a favor and submit one. Without a video, it is very difficult to communicate that the installation will be fun, real, and important. The immediate assumption, right or wrong, is that if you don't have a video, there is a good reason (for example, your work is not implemented yet), and this works against you big time. The jury is composed of people who have done installations and are very forgiving of "bugs." They recognize that the video was made in January and the installation is in July, and that there is lots of time to improve things. So you are always better off showing what you have and what you need to do, than leaving it up to the jury's imagination.

While you want to go out of your way to convince the jury that your installation is real, you should be honest about the risks and remaining work. The jury may not be experts in your field, but they are generally experts in detecting unsubstantiated claims (BS).

Is it important?
One of the objectives of Enhanced Realities is to expose people to novel techniques (or novel applications of known techniques) that have the potential for becoming mainstream several years in the future. This is what I mean by importance. Is there something about the installation that is a harbinger of things to come? In your proposal and video, you should make this very clear. In general, importance should be read as "technical importance" as opposed to "sociological importance." Ideally, your installation should be based on research that is worthy of being presented in SIGGRAPH Papers or Sketches (and in the best of all possible worlds, is being presented in one of those programs), or that you are planning on submitting as a Paper next year. If you believe that the real importance of your installation lies in its sociological implications, don't be deterred, but be sure to present your rationale in a clear and down-to-earth manner.

Jury Statement **Andrew Glassner** Microsoft Research

Even before I started reading the proposals for SIGGRAPH 98's Enhanced Realities, I wanted to accept them all. I knew that every proposal represented the work of one or more people who wanted to share their coolest ideas with the rest of the SIGGRAPH community. Since I like to support imaginative and creative thinking whenever I can, I wanted to invite all the proposers to come along and show their stuff.

But that wouldn't be fair to the SIGGRAPH 98 attendees, whose time at the conference is already in short supply. Our job as a jury was to select those entries that would really reward someone's time and effort. The ideal Enhanced Realities entry would inspire a new idea or cause someone to laugh and grab a friend, insisting that they had to come and check out the piece.

There are at least two ways to think about evaluating pieces. One way is to be selfish and simply judge each entry by its impact on one's own self. Another way is to consciously act as a representative for other viewers and try to choose entries that would delight and excite them. Our chair, Janet McAndless, did a great job of selecting people who represented a cross-section of the SIGGRAPH crowd: artists, scientists, programmers, dreamers, and so on. Janet's up-front work in assembling the committee meant that we were free to be selfish. This meant that as a juror, I could operate as my mind and gut told me to.

Because I believed that every entry was a sincere proposal, I wanted to extract every bit of good stuff from each one, so that I'd have a chance to really evaluate how each work hit me. Unfortunately, many of the proposals didn't give themselves a fair chance. The first big problem was ambiguity. Some proposals were simply too short to explain what they were about.

Others were unclear, either because they were couched in jargon or simply poorly written. We had several proposals that were filled with "art speak." For those who are unfamiliar with this form of expression, it is the intellectual language of art that can obscure as much as it illuminates. Art speak would refer to a pencil as a "trans-physical emotional/intellectual realizer, directly projecting mental thought processes without correlated referents into concrete manifestations of private and public iconography." You get the idea. It is almost impossible to understand what this stuff refers to.

Some projects were simply without purpose, and amazingly, sometimes information was simply missing. We would see pictures or a video of a gadget or a system, and we'd be left wondering why they bothered. What good is this thing? What problem does it solve? What questions does it pose? Why is it different or better than other ways to accomplish the goal? Bottom line: why did the authors do this? When we were sufficiently bewildered, we had to let the proposal go into the reject bin.

A related problem was proposals that stressed the wrong thing. There were some projects that had very cool technology applied to toy problems, when the ultimate purpose was much more interesting and valuable. We were a little more generous to these on the whole, since so much of SIGGRAPH is about technology, and we imagined that most people would see the applications on their own if the technology was sufficiently well presented.

A few of the proposals seemed to try to enhance their credentials by referencing the many fine people and reviewers who had already said good things about the work. If a piece really had some stellar reviews, I didn't mind seeing them mentioned as supporting documentation. However, some proposals dropped names egregiously. I have to admit that I tended to think that any submission that flaunted its credentials was probably doing so to cover something up, and I usually found that something. I would have preferred that the authors let the piece stand on its own merits.

Finally, business blurbs were rejected pretty quickly. We were aware that some of the proposals represented the work of people in commercial enterprises, and that wasn't a problem. But when the submission made it pretty clear that the proposal itself was simply an advertisement for a commercial product, we were rather more careful, since we didn't want Enhanced Realities to turn into a smaller version of the SIGGRAPH 98 Exhibition.

The proposals I liked best were short, clear, well-written and illustrated, and identified the value of the piece. I appreciated proposals that identified what was already working, what could probably be done in time for the conference, and what was just blue-sky. I liked proposals that didn't strain to sell the work, but simply shared why it was done, what was cool about it, and why other people would be turned on to see it. Enthusiasm is always great to see, but the best proposals also shared the lasting value of the work. Many of the proposals that we accepted had their purpose and value stated crisply and clearly up front, and then expanded on that basic idea to show its implications.

I hope that when you tour Enhanced Realities you'll find some presentations that make you laugh, some that make you think, some that give you great new ideas, and some that give you an optimistic picture of the future we're all creating together.

Jury Statement Kathryn Saunders Royal Ontario Museum

Seventeen projects were selected by our jury from the 68 entries originally submitted. This year's selections continue to "raise the bar" of creative and technological excellence that is the hallmark of Enhanced Realities.

While we used the same two-step jury process (screening and finals), each of us on the jury employed our own criteria for grading the submissions. In selecting the jury, the chair astutely combined industry professionals from research, artistic, and production backgrounds, which gave rise to some lively discussions and piercing insights.

I evaluated the projects using the following criteria:

- The level of innovation
- The strength of the idea
- The quality of the content and implementation
- Audience appeal
- Significance and potential influence

Level of Innovation

Certainly the most common question asked by the jury was: "What is the innovative or groundbreaking technology demonstrated here?" Submissions such as Object-Oriented Displays managed to communicate an impressive level of innovation by using a series of diagrams to describe the technology in detail, indicating what the components were and how they were configured. Another project appeared to be based on the same level of innovative research. The jury was very interested and wanted it to be included in the ranks of the winners, but we had to give it a failing grade because it was so poorly documented. It lacked accompanying images, video, or Web site, so we couldn't reasonably take the risk of including it. Ten percent of the submissions were disqualified because of lack of documentation.

Strength of the Idea

There is a saying that computer graphics are the solution to a nonexistent problem. I, however, was completely moved when I came across submissions with strong ideas and real resonance. Stretchable Music with Laser Range Finder is one such project. Participants create music by moving their hands to stretch and distort graphics that in turn control different sounds, rhythms, and patterns. By incorporating an innovative laser hand tracker and superb programming, the effect is seamless, intuitive, and very satisfying. Other ideas were strong partly because they were so "simple." inTouch, for example, did not demonstrate breakthrough technology. Rather, its strength was its unabashed response to a basic human need for tactile communication between two people separated by distance.

Quality of Content and Implementation

I looked for content that was not only engaging and beautiful, but offered something beyond the expected. For example, Media & Mythology uses role-playing within a multi-user domain as a vehicle to learn about ancient mythology. When the jury met, much of Media & Mythology was not yet produced. However, using images of past work and conceptual sketches, the submitters were able to convince us that they had the talent and the resources to pull it off. To be accepted, contributors had to reasonably demonstrate that they would be able to control or direct a unified, consistent, holistic experience, which is not easy for most people. Entrants who worked as part of a collaborative team ranging from research scientists to artists, to shape the new technologies into meaningful experiences, were more successful. Those who actively solicited feedback from the Enhanced Realities chair enjoyed even more success.

Audience Appeal

Each of the successful submissions proposed experiences that allowed people to actively participate. Although the type of interaction varied, the jury was completely aligned on the fact that the experiences had to be fun and engaging. From the immediately accessible interactive environment of the HoloWall and Mass Hallucination to the more process-driven experience of Shall We Dance? and Virtual Head, the successful projects are to be commended overall for their interactive focus.

Significance and Potential Influence

Enhanced Realities showcases breakthrough technologies combined with innovative content. Next year, the industry will have advanced through the evolution or deviation of these and other ideas. Submissions that have the potential to act as a catalyst or influence future technologies were favored by the jury. Although technological breakthroughs received the highest marks, incremental improvements and innovations were also accepted.
One project, which resurfaced from last year, demonstrated significant advancements in its innovation. Although the purpose of the project remained constant, the technologies used in 1998 were a marked improvement over the 1997 submission. The contributor wrote that during SIGGRAPH 97, he met someone who suggested another way of approaching the problem. I liked that. It reinforced SIGGRAPH's role in the exchange and development of ideas that are actively shaping our future.

HoloWall: Interactive Digital Surfaces

HoloWall is an interactive wall system that allows visitors to interact with digital information displayed on the wall surface without using any special pointing devices. The combination of infrared cameras and infrared lights installed behind the wall enables recognition of human bodies, hands, or any other physical objects that are close enough to the wall surface. Visitors can use both hands simultaneously. Body shape can also be an input to the system. HoloWall demonstrates several interactive environments, including a world of autonomous digital insects that respond to body movements and an interactive sound environment that reactively creates music sequences based on the user's actions.

HoloWall demonstration at NTT ICC.
Courtesy of NTT InterCommunication Center [ICC]

Jun Rekimoto
Sony Computer Science Laboratory, Inc.
3-14-13 Higashigotanda, Shinagawa-ku
Tokyo 141-0022 Japan
rekimoto@csl.sony.co.jp
www.csl.sony.co.jp/person/rekimoto/holowall/

Masaaki Oka
Nobuyuki Matsushita
Sony Corporation

Hideki Koike
University of Electro-Communications

Collaborators
Hiroaki Tobita
Jun Nagahara
Tota Hasegawa
Hidenori Kurasawa
Hiroyuki Sueno

Swamped! Using Plush Toys to Direct Autonomous Animated Characters

Synthetic Characters Group
& Vision and Modeling Group
MIT Media Lab

Swamped! is a multi-user interactive environment in which instrumented plush toys are used as an iconic and tangible interface to influence autonomous animated characters. Each character has a distinct personality and decides in real time what it should do based on its perception of its environment, its motivational and emotional state, and input from its "conscience," the guest. A guest can influence how a given character acts and feels by manipulating a stuffed animal corresponding to the character. For example, the guest could direct her character's attention by moving the stuffed animal's head, comfort it by stroking its belly, or have it wave at another character by waving its arm. Automatic camera control is used to help reveal the emotional content of each scene. By combining research in autonomous character design, automatic camera control, tangible interfaces and action interpretation, Swamped! seeks to create a rich, evocative and novel experience.

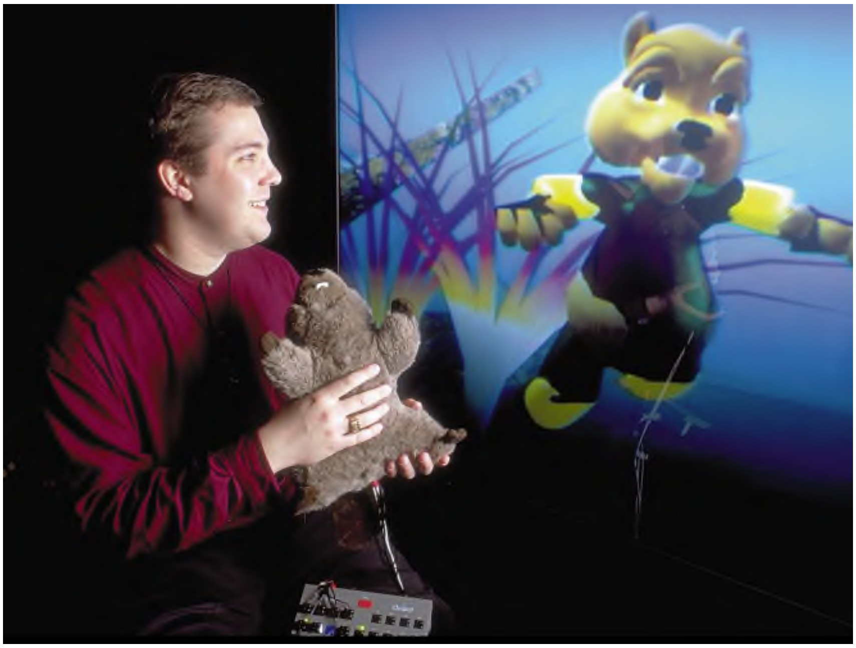

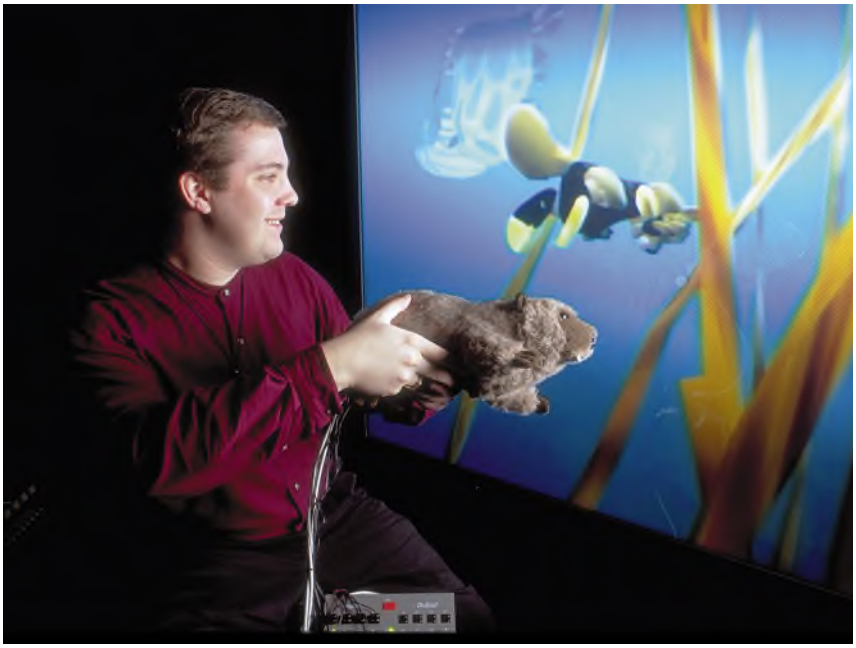

Bruce M. Blumberg
Synthetic Characters Group
MIT Media Lab
E15-311, 20 Ames Street
Cambridge, Massachusetts 02139 USA
bruce@media.mit.edu
characters.www.media.mit.edu/groups/characters/

Collaborators
Bruce Blumberg
Michael Patrick Johnson
Michal Hlavac
Christopher Kline
Kenneth Russell
Bill Tomlinson
Song-Yee Yoon
Agnieszka Meyro (Administrative Assistant)
Synthetic Characters Group, MIT Media Lab

Aaron Bobick
Andrew Wilson
Vision and Modeling Group, MIT Media Lab

And the following undergraduates:
Jed Wahl
Jeremy Lueck
Todd Nightingale
Javorka Saracevic
Damian Isla
Zoe Chelsea Teegarden

AR² Hockey

Toshikazu Ohshima, Kiyohide Satoh, Hiroyuke Yamamoto, and Hideyuki Tamura
Mixed Reality Systems Laboratory, Inc.

In AR² Hockey (Augmented Reality AiR Hockey), players share a physical game field, mallets, and a virtual puck to play air hockey in simultaneously shared physical and virtual space. They can also communicate with each other through the mixed space. Since real-time, accurate registration between both spaces and players is crucial to playing the game, a video-rate registration algorithm is implemented with commercial head-trackers and video cameras attached to optical see-through head-mounted displays. Our collaborative AR system achieves higher interactivity than a totally immersive VR system.

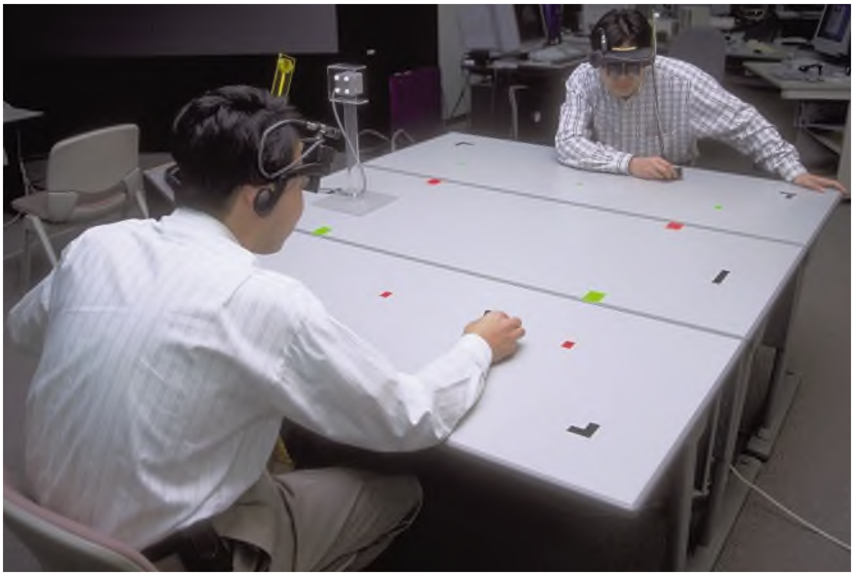

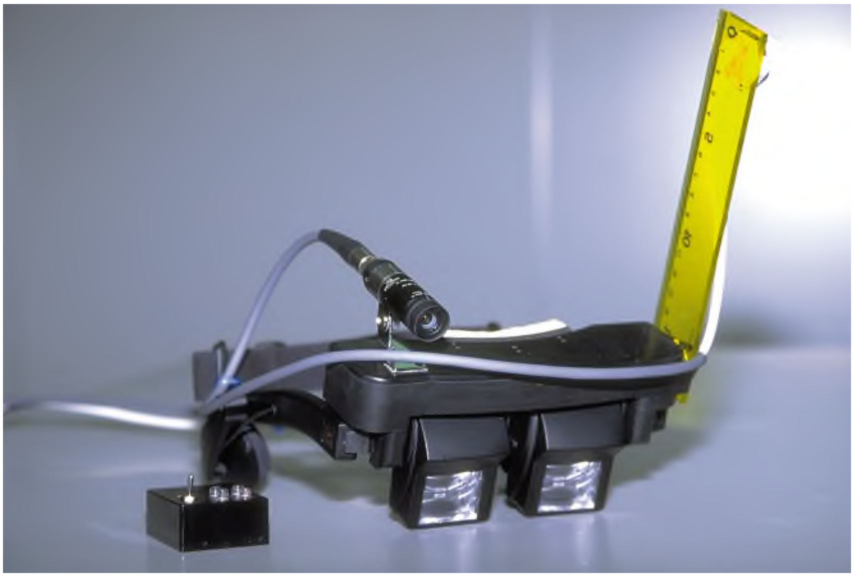

Toshikazu Ohshima
Mixed Reality Systems Laboratory, Inc.
6-145 Hanasaki-cho
Nishi-ku, Yokohama 220-0022 Japan
ohshima@mr-system.com
www.mr-system.com/

Collaborators
Toshikazu Ohshima
Kiyohide Satoh
Hiroyuke Yamamoto
Hideyuki Tamura

PingPongPlus

Tangible Media Group
MIT Media Lab

PingPongPlus is a digitally enhanced version of the classic ping-pong game. Various audio and visual augmentations have been added to a conventional ping-pong table with a non-invasive, sound-based ball tracking system. The "reactive table" displays patterns of light and shadow as a game is played, and the rhythm and style of play drives accompanying sound. At times, the game is subtly enhanced, and sometimes it is powerfully changed. In one mode, the table appears to be covered with water, so that playing on it creates patterns of subtle ripples. In another mode, images that race around the table change the entire scoring system and method of play. The goal of the project is to explore systems for collaborative play that push the physical world back into the forefront of design, without relying on simple GUI controllers, such as a mouse, keyboard, and joystick.

Craig Wisneski
MIT Media Lab
Massachusetts Institute of Technology
E15-452, 20 Ames Street
Cambridge, Massachusetts 02139 USA
wiz@media.mit.edu
tangible.media.mit.edu/projects/pingpongplus.html

Collaborators
Craig Wisneski
Julian Orbanes
Hiroshi Ishii

Object-Oriented Displays

Naoki Kawakami, Masahiko Inami, Yasuyuki Yanagida, and Susumu Tachi
The University of Tokyo

In Object-Oriented Displays, users perceive and operate a virtual object as if it were real. Design and implementation of three types of object-oriented displays are demonstrated: MEDIA-Ace, a liquid crystal display (LCD) and position sensor; MEDIA-Cube, a position sensor and four LCDs arranged in the shape of a cubic body; and MEDIA-Crystal, which uses optical projection.

Naoki Kawakami, Masahiko Inami, Yasuyuki Yanagida, and Susumu Tachi
Tachi Lab. MEIP, The Faculty of Engineering
The University of Tokyo
7-3-1 Hongo Bunkyo-ku
Tokyo 113-8656 Japan
kawakami@star.t.u-tokyo.ac.jp

Collaborators
Naoki Kawakami
Masahiko Inami
Yasuyuki Yanagida
Susumu Tachi

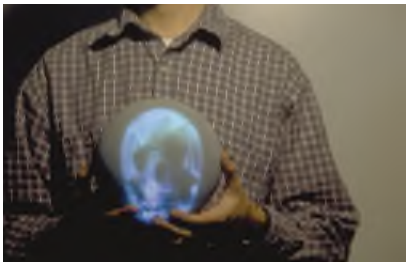

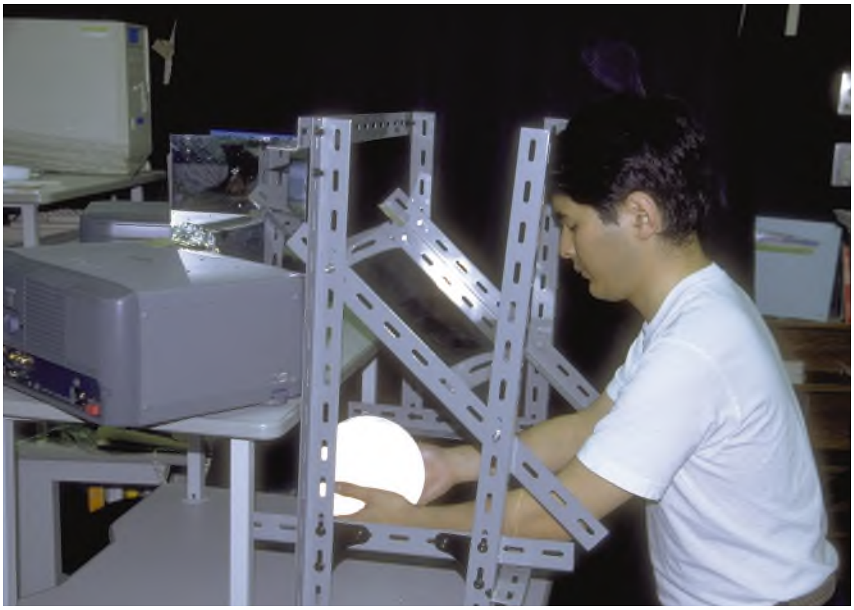

Mass Hallucination

T. Darrell, M. Harville, G. Gordon, J. Woodfill
Interval Research Corporation

This imaging display changes according to the number of people watching it, their behaviors, and whether they've watched the device before. It is reflexive: the displayed image is a function of the people watching the display. It encourages crowds of people to collectively manipulate the display with their bodies or faces. Yet it is also personal, in that it can recognize the appearance of a user for short-to-medium periods of time and tailor the display accordingly. As in Magic Morphin' Mirror, a SIGGRAPH 97 Electric Garden project by the same group, this display captures video along the same optical axis as video is displayed, so images of observers can be directly manipulated, composited, or distorted on the display. In contrast to the previous work, which only considered a single user at a time and had no persistence after they left, this display is designed to visually track a crowd of people and provide a shared graphical experience. It also tracks users over time through multiple sessions. We show that continuity/consistency of experience across multiple simultaneous users, or a single user at a time, is possible.

Trevor Darrell
Interval Research Corporation
1801 Page Mill Road, Building C
Palo Alto, California 94304 USA
trevor@interval.com

Foot Interface: Fantastic Phantom Slipper

Akihiko Shirai, Masaru Sato, Yuichiro Kume, and Machiko Kusahara

People should be able to use their feet just as freely in a virtual environment as they do in the real world. Wearable interfaces should not cause psychological and/or physical discomforts. This slipper-like multi-modal interface is based on those two assumptions. It features a slipper interface with cyberworlds. Each foot's movement is measured in real time with an optical motion capture system, and feedback signals are transmitted to the soles. Phantom sensations elicited by multiple tactile stimuli allow transmission of complicated feedback information such as objects moving around the feet. Optical markers for motion capture and vibrators for tactile stimulation are installed in the slippers. Players interact with virtual objects projected onto a floor screen, sense them, and use them to play games. The system runs on a single PC.

Yuichiro Kume
Tokyo Institute of Polytechnics
1583 Iiyama,
Atsugi, Kanagawa 243-0297 Japan
kume@photo.t-kougei.ac.jp
laplace.photo.t-kougei.ac.jp/

Collaborators
Akihiko Shirai
Masaru Sato
Yuichiro Kume
Machiko Kusahara

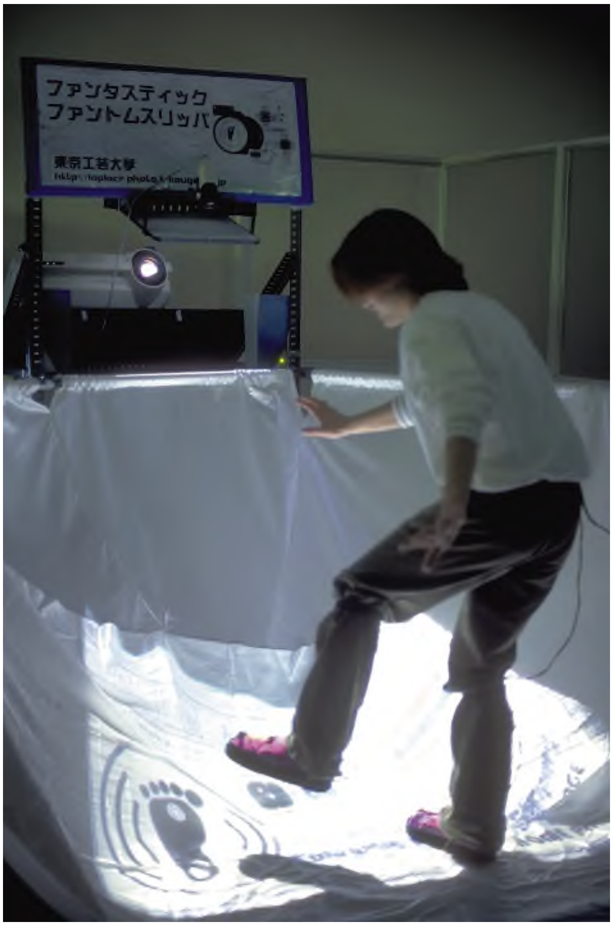

inTouch

Tangible Media Group
MIT Media Lab

Touch is a fundamental aspect of interpersonal communication. Yet while many traditional technologies allow communication through sound or image, none is designed for expression through touch. The goal of inTouch is to bridge this gap by creating a physical link between users separated by distance. InTouch consists of two separate identical objects, each consisting of three cylindrical rollers mounted on a base. The two objects behave as if corresponding rollers are physically connected, but in reality, the objects are only virtually linked. Sensors are used to monitor the states of the rollers, and computer-controlled motors synchronize those states, creating the illusion that distant users are interacting through a single, shared physical object.

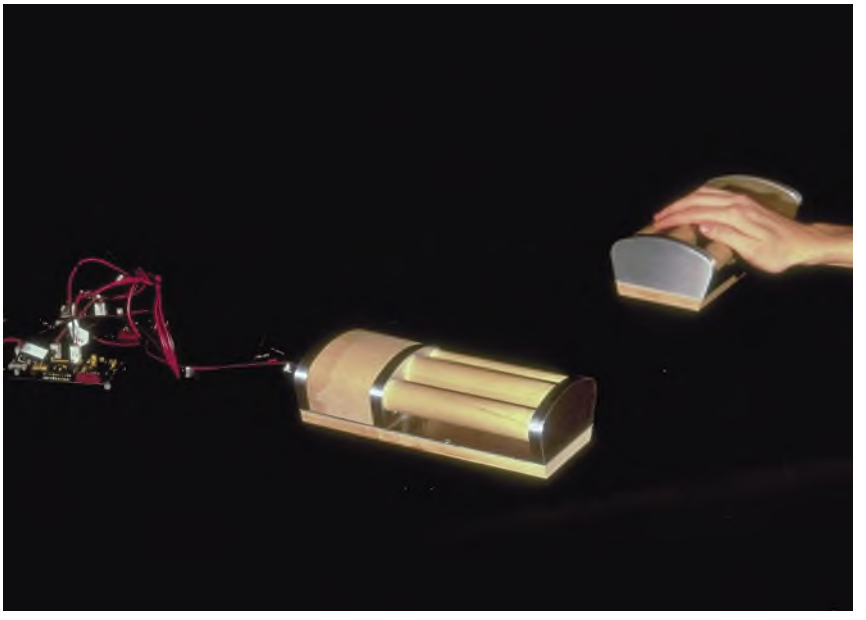

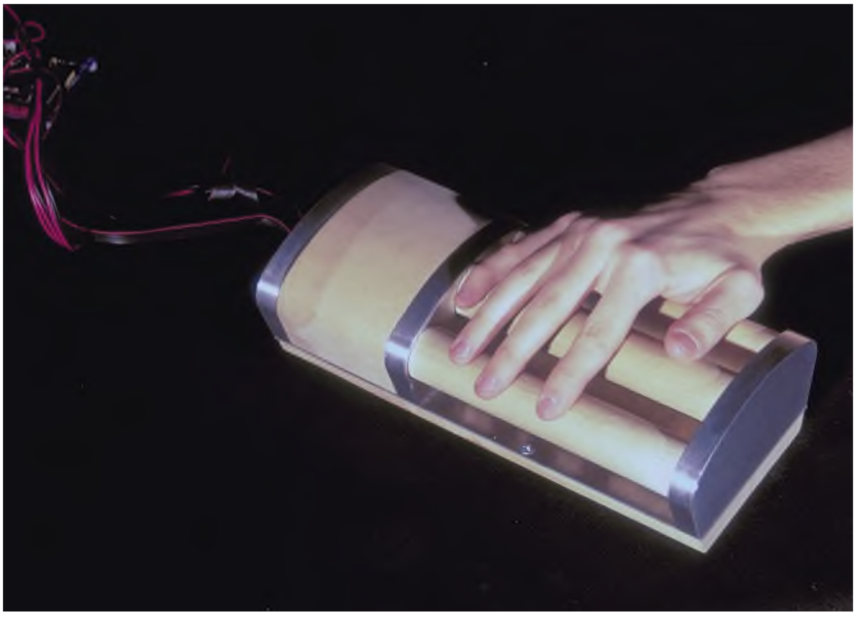

Scott Brave
MIT Media Lab
Massachusetts Institute of Technology
E15-468C, 20 Ames Street
Cambridge, Massachusetts 02139 USA
brave@media.mit.edu
tangible.media.mit.edu/projects/intouch.html

Collaborators
Scott Brave
Andrew Dahley
Phil Frei
Hiroshi Ishii

Virtual FishTank

Nearlife, Inc.

The Virtual FishTank is a simulated aquatic environment featuring a 400-square-foot tank populated by whimsical and dynamic fish.

Participants can:

- Create their own fish.
- Design behaviors for their fish.
- Observe their fish interacting with other fish.
- Manipulate behavioral rules for a group of fish.
- Discover how these behaviors can emulate schooling.
- Analyze emerging patterns.

Through real-time 3D graphics, visitors are introduced to ideas from the sciences of complexity – ideas that explain not only ecosystems, but also economic markets, immune systems, and traffic jams. In particular, visitors learn how complex patterns arise from simple rules. The first version of Virtual FishTank opens at The Computer Museum in Boston in June 1998. A second version will travel nationally to other science museums and aquariums.

Stacy Koumbis
147 Sherman Street
Cambridge, Massachusetts 02140 USA
stacy@nearlife.com
www.nearlife.com

Collaborators
Tinsley Galyean
Sheri Galyean
Brian Knep
Henry Kaufman
Stacy Koumbis
Aubrey Francois
George Bird
David Zung
David Friend
Karen Sideman
Patrick Porter
Scott Yu
Darrin Bascome
Mitchel Resnick and the MIT Media Lab
Oliver Strimpel and The Computer Museum

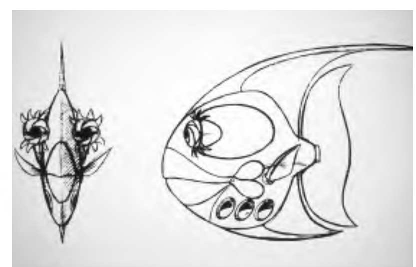

Haptic Screen

Hiroo Iwata
University of Tsukuba

Haptic Screen is a new force-feedback device that deforms itself to present shapes of virtual objects. Typical force-feedback devices use a grip or thimble, but users of Haptic Screen can touch the virtual object without wearing anything. Haptic Screen employs an elastic surface made of rubber. A 6 X 6 array of 36 actuators deforms the surface and controls its hardness according to the force applied by the user. An image of the virtual object is projected onto the elastic surface so that the user can directly touch the image and feel its rigidity.

Hiroo Iwata
Institute of Engineering Mechanics
University of Tsukuba
Tsukuba, 305 Japan
iwata@kz.tsukuba.ac.jp
intron.kz.tsukuba.ac.jp

Collaborators
Hiroo Iwata
Hironori Nakagawa
Wataru Hashimoto
Fumitaka Nakaizumi
Iwata Laboratory
Institute of Engineering Mechanics

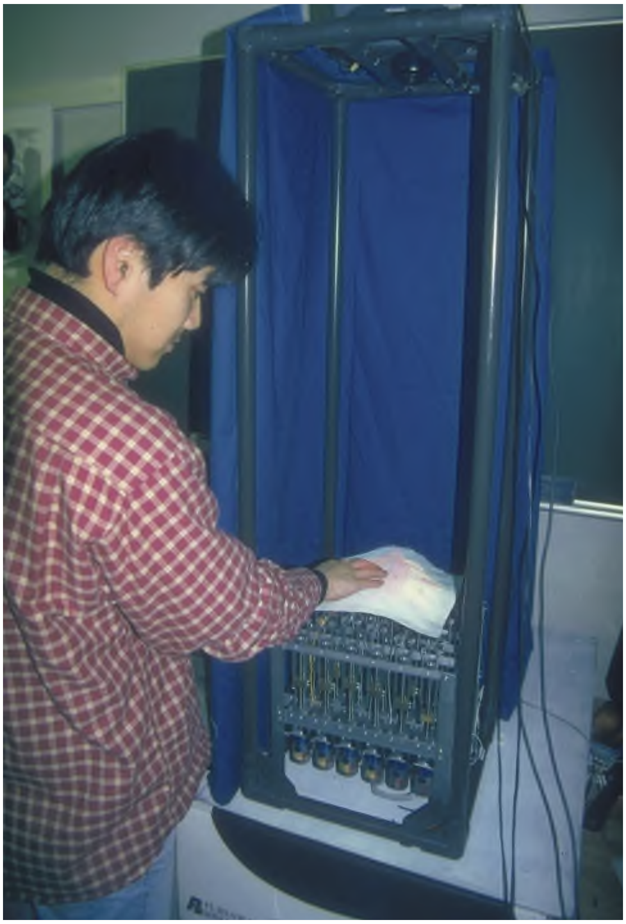

Natural 3D Display System Using Holographic Optical Element

In this natural 3D display system, a holographic optical element (HOE) overcomes conflicts between convergence and accommodation. Users experience clear stereoscopic vision, without glasses, of a broad field of view. With its multiple-focus HOE, the system offers two pairs of viewing points in back-and-forth or horizontal locations.

Koji Yamasaki
Laboratories of Image Information
Science and Technology
1-1-8-3F, Shinsenri-Nishi, Toyonaka,
Osaka 565-0083 Japan
yamasaki@senri.image-lab.or.jp

Collaborators
Koji Yamasaki and Masaaki Okamoto
Laboratories of Image Information
Science and Technology

Naoki Nakanishi
Noritsu Koki Co., Ltd.

Eiji Shimizu
Osaka City University

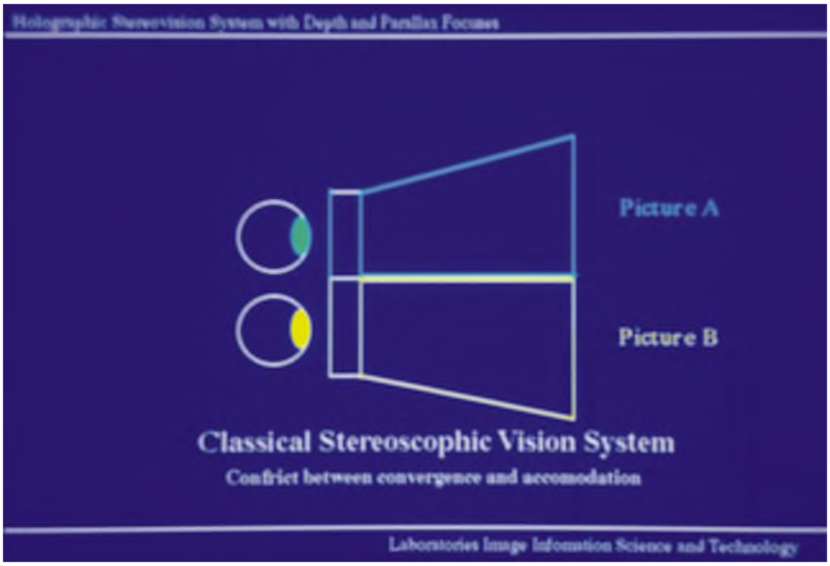

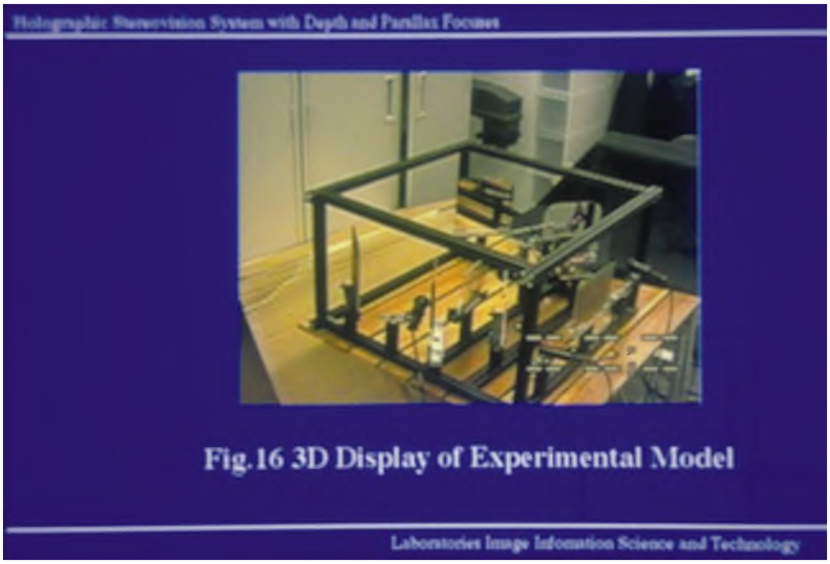

Direct Watch & Touch

Laboratories of Image Information Science and Technology

This 3D display offers access to a virtual stereoscopic world without special glasses. When users “touch” the world with real tools (for example, a hammer, a surgical knife, a wrench, tweezers, etc.), directly and interactively, they hear and feel contact and transform virtual objects. This binocular parallax display combines virtual and real environments in full, high-resolution (XGA) color. It is a new approach to virtual reality that handles virtual objects with “real” tactile feedback.

Takahisa Ando
Laboratories of Image Information Science and Technology
Daiichi-Kasai Senri-Chuo Bldg. 3F, 1-1-8,
Shinsenri-Nishimachi, Toyonaka,
Osaka 565 Japan
ando@image-lab.or.jp
www.image-lab.or.jp/

Collaborators
Takahisa Ando
Eiji Shimizu
Makoto Sato
Goro Hamagishi
Atsushi Suwa

Media & Mythology

Institute for Simulation & Training

In ancient times, mythology was the high-tech method for storing data on a society's history, rituals, and ethical systems. The paradigm in use for these early information systems was storytelling. Media & Mythology explores the link between traditional mythologies from several cultures and new technology/new media. Man and Minotaur allows visitors a chance to portray the two ancient combatants and the gods that taunt them within a fully immersive, synthetic version of Dedalus' Labyrinth in ancient Crete. In Video Totem, expressionistic visitors create and view their own mythologies on a large digital totem pole. Dear Oracle integrates contemporary media into traditional soothsaying. The result is a new form of oracle: digital divination.

Producer
Kimberly Abel Parsons
Visual Systems Laboratory
Institute for Simulation & Training
3280 Progress Drive
Orlando, Florida 32826 USA
kparsons@ist.ucf.edu

Creative Director
James B. Parsons

Art Director
Anthony B. Mickle

Contributing Artists and Scientists
Visual Systems Laboratory
Institute for Simulation & Training
Students of the CREAT Program
University of Central Florida

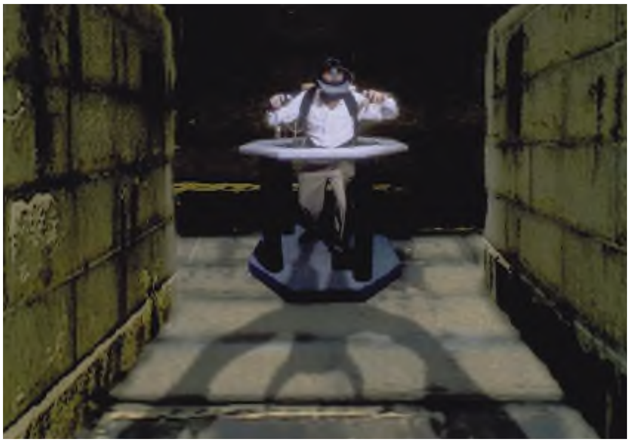

Natural Pointing Techniques Using a Finger-Mounted Direct Pointing Device

John Sibert and Mehmet Gokturk
The George Washington University

Pointing with the index finger is a natural way to select an object, and if it can be incorporated into human-computer interaction technology, a significant benefit will be obtained for certain applications. This demonstration presents a prototype solution.

Based on an infrared signal power density weighing principle, a small infrared emitter on the user's finger and multiple receivers placed around the laptop screen generate data for a low-cost microprocessor system. The microprocessor sends its output to a laptop computer, where it is used to determine coordinates for the cursor location. The prototype is not only a proof of concept. It is also a tool for further research on human performance in pointing and further development of interactive techniques.

Wearing the finger-mounted emitter. (The black box on the wrist contains batteries and the modulator circuit.)

John Sibert
Department of Electrical Engineering and Computer Science
The George Washington University
Washington D.C. 20052 USA
sibert@seas.gwu.edu

Collaborators
John Sibert
Mehmet Gokturk
Robert Lindeman
Sang Yoon Lee

The prototype in use.

Virtual Head

Echtzeit GmbH

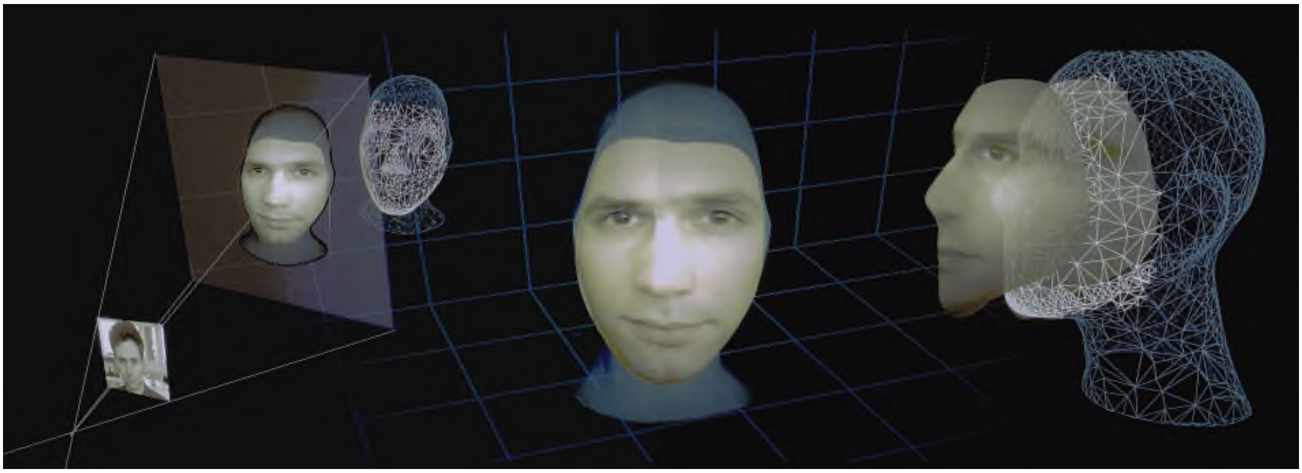

Virtual Head is a new approach that enhances communication in virtual environments and telepresence. It tackles one of the key problems in the field of innovative telecommunication technology: how to represent oneself in virtual environments in such a way that an emotional and natural way of communicating with others is possible?

The Virtual Head conferencing prototype renders three-dimensional images of every communication partner in real-time. It establishes eye-to-eye contact among the communication partners by projecting live-video textures onto 3D geometry of a head. The application translates the head movement so that video-images show the original movements. Compressed video and audio information is exchanged via a high-bandwidth network to establish a remote conferencing scenario. Video and audio are decompressed on both sides, and the images are projected onto a screen.

This approach uses original face images with all their facial expressions and tries to transport the main factors of human communication such as line of gaze, which indicates attention and significantly drives a conversation. According to psychologists, most of the information we remember after talking with somebody is non-verbal. Improving technologies for visual communication that includes a more "emotional" way of meeting each other in virtual environments will become possible with high-bandwidth networks in the very near future.

Thom Brenner
Echtzeit GmbH
Kanstrasse 165
10623 Berlin, Germany
tbrenner@echtzeit.de

Collaborators
Thom Brenner
Henrik Battke
Ilja Radusch
Raimo Ihle

Stretchable Music with Laser Range Finder

Stretchable Music with Laser Range Finder combines an innovative, graphical, interactive music system with a state-of-the-art laser tracking device. An abstract graphical representation of a musical piece is projected onto a large vertical display surface. Users are invited to shape musical layers by pulling and stretching animated objects with natural, unencumbered hand movements. Each of the graphical objects is specifically designed to represent and control a particular bit of musical content. Objects incorporate simple behaviors and simulated physical properties to generate unique sonic personalities that contribute to their overall musical aesthetic. The project uses a scanning laser rangefinder to track multiple hands in a plane just forward of the projection surface. Using quadrature-phase detection, this inexpensive device can locate up to six independent points in a plane with cm-scale accuracy at up to 30 Hz. Bare hands can be tracked without sensitivity to background light and complexion to within a four-meter radius.

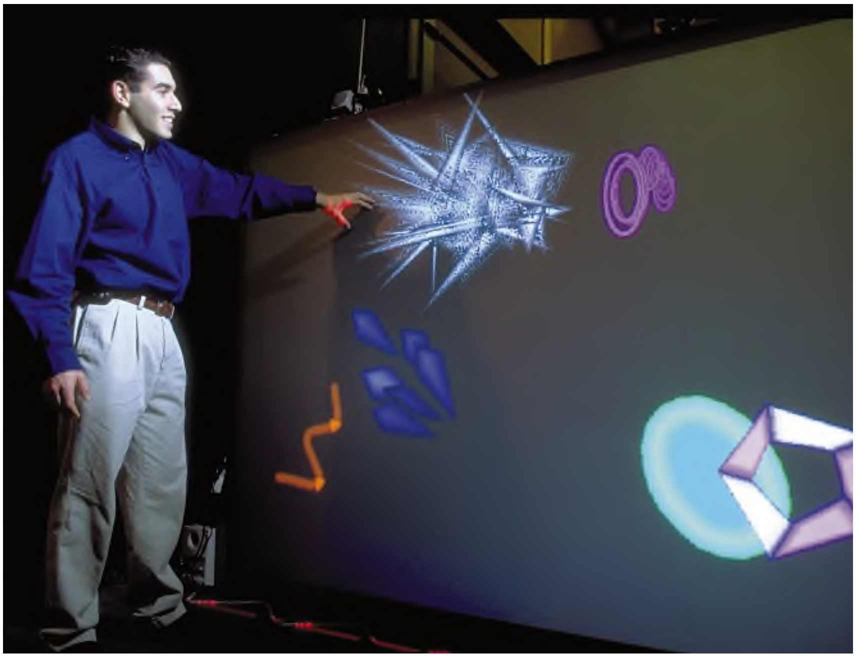

Pete Rice and Joshua Strickon
MIT Media Lab
Massachusetts Institute of Technology
E15-495, 20 Ames Street
Cambridge, Massachusetts 02139 USA
strickon@media.mit.edu
brainop.media.mit.edu/~strickon/siggraph.html

Collaborator
Joe Paradiso

Shall We Dance?

ATR Media Integration & Communication Research Lab and University of Maryland

Real-time 3D computer vision gives users control over both the movement and facial expression of a virtual puppet and the music to which the puppet "dances." Multiple cameras observe a person, and human silhouette analysis achieves real-time 3D estimation of human postures. Facial expressions are estimated from images acquired by a viewing-direction controllable camera, so that the face can be tracked. From the facial images, deformations of each facial component are estimated. The estimated body postures and facial expressions are reproduced in the puppet model by deforming the model according to the estimated data. All the estimation and rendering processes run in real time on PC-based systems. Attendees can see themselves dancing in a virtual scene as virtual puppets.

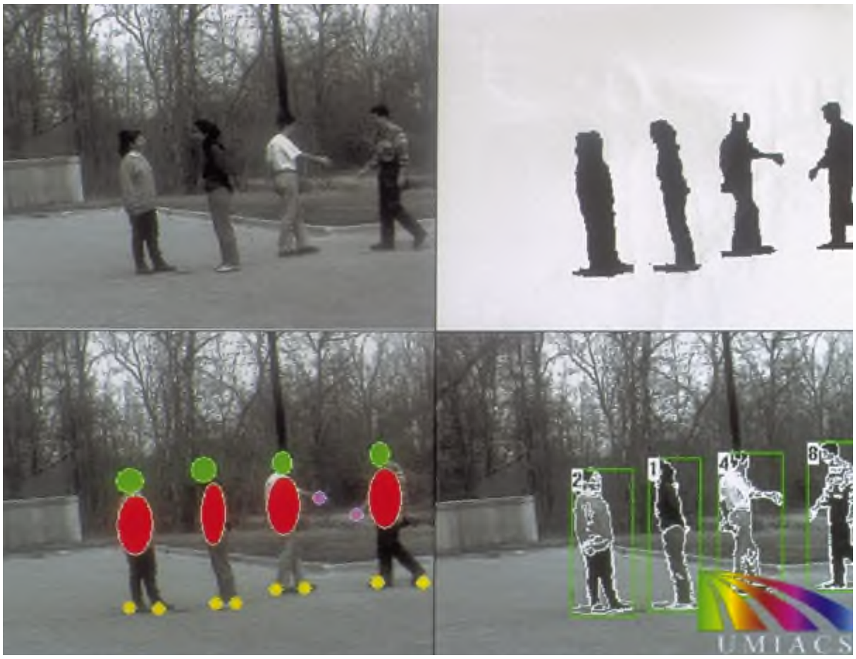

Kazuyuki Ebihara
ATR Media Integration & Communication Research Lab
2-2 Hikaridai Seika-cho Soraku-gun
Kyoto 631 Japan
ebihara@mic.atr.co.jp

Collaborators
Kazuyuki Ebihara
Jun Kurumisawa
Tatsumi Sakaguchi
Jun Ohya
ATR Media Integration & Communication Research Lab

Larry S. Davis
Thanarat Horprasert
R. Ismail Haritaoglu
University of Maryland

Digital Pavilions

Introduction

"I used to think that cyberspace was 50 years away. What I thought was 50 years away, was only 10 years away. And what I thought was 10 years away… it was already here. I just wasn't aware of it yet."

Bruce Sterling

When you set out to create a collective vision of even near-future technology, you learn some surprising things along the way. Digital Pavilions is the result of much behind-the-scenes effort and dedication by technology visionaries whose goal is to demonstrate better living through networks and computer graphics.We hope to inspire the technical and artistic cultures within the SIGGRAPH community to reconsider our impact on new media, to dream what may be possible for a "personal reality," and to rediscover our interconnectedness in a global society. What's in this future forever changed by instant, increasingly high-bandwidth contact across the globe?

Digital Pavilions offers an imaginative look at new answers on the horizon, and at current research on near-future collaborative technologies that will connect us ever more closely to our CPUs, our families, and our coworkers. With the latest advances in many emerging areas, from natural language processing to gesture recognition, we will soon be able to interact with the digital world as naturally as we do with our neighbors in everyday life. Families separated by hundreds or thousands of miles can easily converse and visually share experiences using hand-held graphical devices that you might find on sale in consumer electronics stores this holiday season.

Through the integrated Digital Pavilions immersive reality installation ("the living room of the future") new experiences, including travel to far-off destinations like the Alaskan wilderness or attending SIGGRAPH 99 in Los Angeles, become possible for the SIGGRAPH 98 audience in Orlando! And forget armchair quarterbacking. Now you can sit on the sidelines directing your own sportscast via an interactive interface to ESPN's Winter-X Games coverage.

Digital Pavilions envisions the convergence of leading-edge technologies that demonstrate near-future ways to communicate, collaborate, and interact. We invite you to share in our collective vision through a showcase of applications, installations, and sessions as we look toward our collaborative future in an increasingly connected world.

Omar Ahmad
Chair, Digital Pavilions

Netscape Communications Corp.
501 East Middlefield Road
Mountain View, California 94043 USA
Omar_Ahmad@siggraph.org

Digital Pavilions Committee

Rob Lewis
MultiGen, Inc.

Paul Lipsky
Manhattan Transfer

Janet McAndless
Sony Pictures Imageworks

Joe Munkeby
LWUP

Islands of Adventure in Cyberspace

**University of Central Florida,
Universal Studios, and Real 3D**

Islands of Adventure is a series of imaginary experiences based on a real place-to-be: the Universal Studios Islands of Adventure theme park in Orlando, now under construction. Working with Universal Creative and Real3D, the University of Central Florida's CREAT Digital Media Program, a student-faculty team, is pushing the edges of what is possible in VRML by developing and rendering elaborate computer-generated scenery and hand-drawn animation.

The project simulates what a typical Web interaction might be like when homes are routinely equipped with high-performance 3D graphics-capable and sound-capable information systems. SIGGRAPH 98 attendees use Silicon Graphics workstations and high-performance PC equipment equipped with Real3D graphics systems to experience a virtual theme park. Others visit the theme park via a VRML-based Web site.

Contact
J. Michael Moshell
University of Central Florida
Computer Science Department, Box 25000
Orlando, Florida 32816 USA
moshell@cs.ucf.edu
www.creat.cas.ucf.edu/

Collaborators
J. Michael Moshell
David N. Haxton
Dan Silverglate
Charles E. Hughes

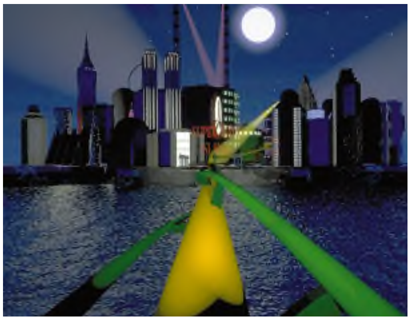

University of Illinois at Chicago

Two roving ImmersaDesks travel throughout SIGGRAPH 98 to:

- Demonstrate networked virtual environments in the Orange County Convention Center
- Run “multi-way” networked environments connected to users at remote sites.
- Showcase stand-alone, interactive, VR projects.

The virtual worlds in Guerilla VR present participatory narratives; push technical limitations to create lush, almost tangible imagery; and launch users into networked interaction with other people.

Among other applications, Guerilla VR features:

- Asteroid A-612, an application that allows several remotely located participants to simultaneously explore a very small planet. It is designed to help teach children that the world is round by allowing them to explore a world that is much more obviously spherical than our own.
- Dimension World, which uses the three dimensions of VR to teach about hypercubes and 4D math. The instructor is remotely located.
- The Thing Growing, where a participant in one location interacts with a virtual character while, in another location, another participant influences the character’s behavior and the progression of the virtual story.

The Electronic Visualization Laboratory’s goals are to demonstrate the feasibility of flexible, accessible VR; increase the number and kinds of people who can be exposed to this medium; and demonstrate its viability and vitality in an ever-increasing variety of institutions, public areas, and private spaces. We are also interested in demonstrating the possibilities of using the next generation of the Internet to create collaborative virtual spaces for work and play.

Contact
Josephine Anstey
Electronic Visualization Laboratory
University of Illinois at Chicago
851 South Morgan Street, Room 1120
Chicago, Illinois 60607 USA
anstey@evl.uic.edu
www.evl.uic.edu/anstey

Collaborators
Joseph Alexander
Josephine Anstey
Hisham Bizri
Tom Coffin
Jim Costigan
Alan Cruz
Greg Dawe
MDolinsky
Tomoko Imai
Ka-Leung Jark
Andy Johnson
Alan Millman
Dave Pape
Maria Roussos
Chris Vasilakis

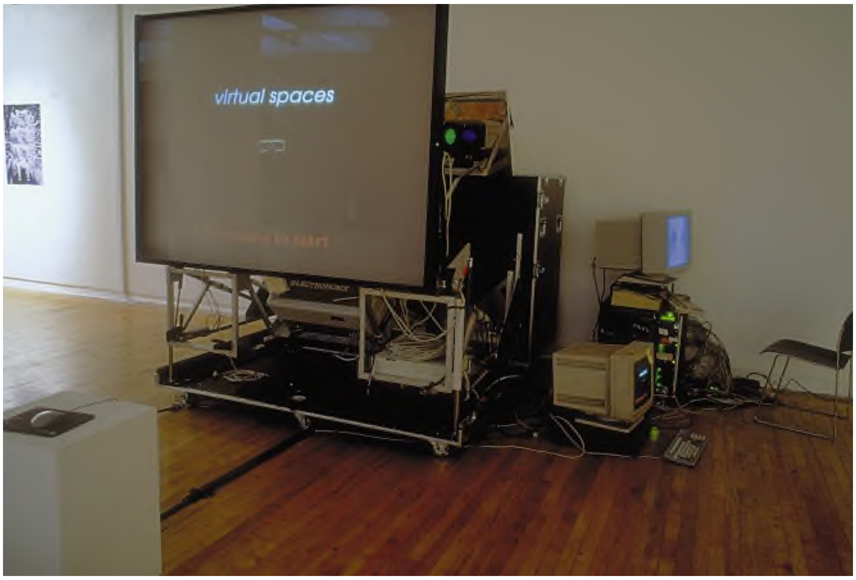

ImmersaDesk2 - a portable, “one box”, user-friendly, VR system.

The Thing Growing
The user is subjected to a virtual character, its whims, and the strange customs of its world.

Mitologies
A narrative environment in which the user takes a strange journey through a rich and intricate labyrinth.

Strait Dope
A figurative world of surreal, visual metaphors and subversive interactions.

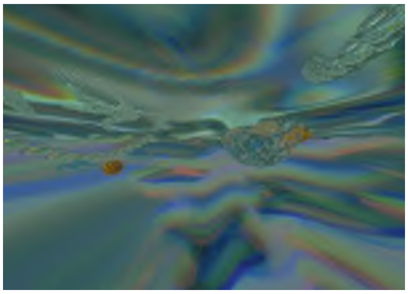

Inner Contact
An abstract environment that uses algorithms to create geometric structures and their behaviors.

GRID
A cyberspace gallery of sculptural mandalas.

Las Meninas
An educational meditation on the Velasquez painting, that plays with the issues of boundaries and reflection.

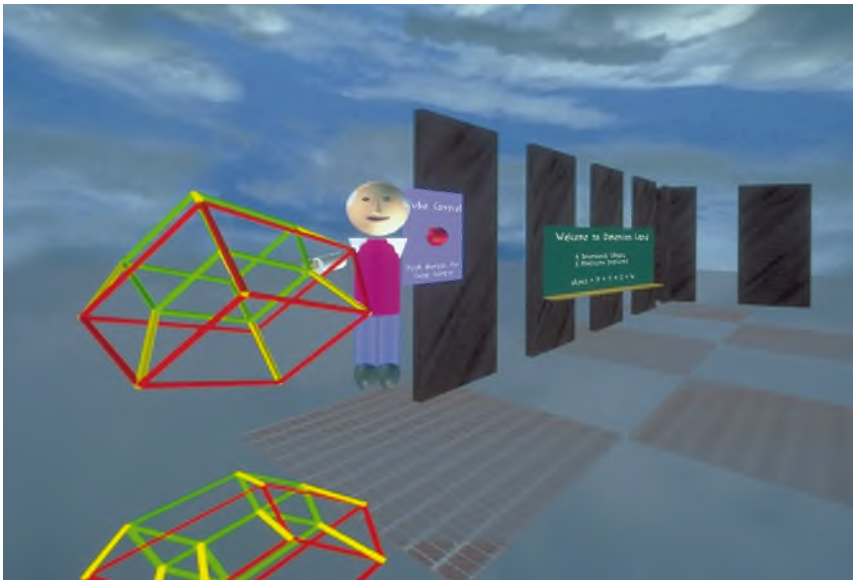

Dimension World uses the three dimensions of VR to teach about hypercubes and 4D math.

Gesture VR: Gesture Interface to Spatial Reality

Bell Laboratories

In this demonstration, users interact with spatial simulations by means of a novel hand gesture recognition interface technology developed at Bell Labs. A freely moving, gloveless hand is the sole input device. Image sequences of the user's hand motions, acquired by video cameras, are processed by a computer program that recognizes gestures and calculates the hand's parameters. This information is used for precise control of navigation in 3D space, for grasping and moving objects on the screen, or to provide a new kind of interface in video games. Users fly through the Yosemite Valley by pointing in 3D; they "grasp" and move objects to compose 3D scenes; they play a video game (such as Doom) in which a character is directed by hand pointing, and game actions are triggered by gestures.

This interactive experience has a very natural feel. GestureVR makes control of complex actions in 3D space very intuitive and simple. Its natural feel and spontaneity are supported by the unique technical qualities of its response time, precision, and robustness.

Contact
Jakub Segen
Bell Laboratories
Room 4E-632
101 Crawfords Corner Road
Holmdel, New Jersey 07733-3030 USA
segen@lucent.com
www.lucent.com/ideas2/innovations/innovators/ti-segen.html

Collaborator
Senthil Kumar
Bell Laboratories

underscore

MIT Media Lab

Hearing music: we know this sort of thing well. Looking at music: there is such a thing as that, but we know it not nearly so well. The one is effortless, you say, egalitarian, accessible to everyone, easy as Play and Pause and Volume. The other is difficult, you assert, specialized, intended for men of bowties and repute, fraught with Staff and Clef and Motif.

But if, you then muse: if there were some kind of artificial bowtie that you could apply, one that led you through a printed score in synchrony with its aural performance and in a way that always made clear the relationship of the seen to the heard...

If, you ponder: if it let you navigate this music, not just with PLAY and PAUSE and REW and FF but also by swooping and diving, peering forward and back, piloting through the score as if above a landscape of notes...

If, you mull: if it gave you breadcrumb-like markers to drop so you could find your way back to the best parts, if it let you create excerpts, if it ran on SGI hardware, if it above all venerated the beauty of musical typography...

Well, then...

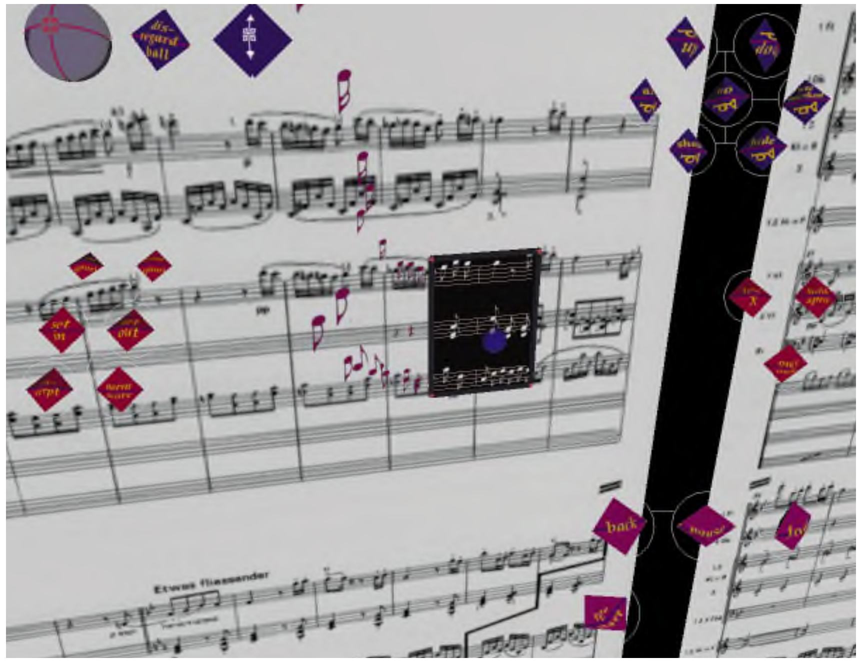

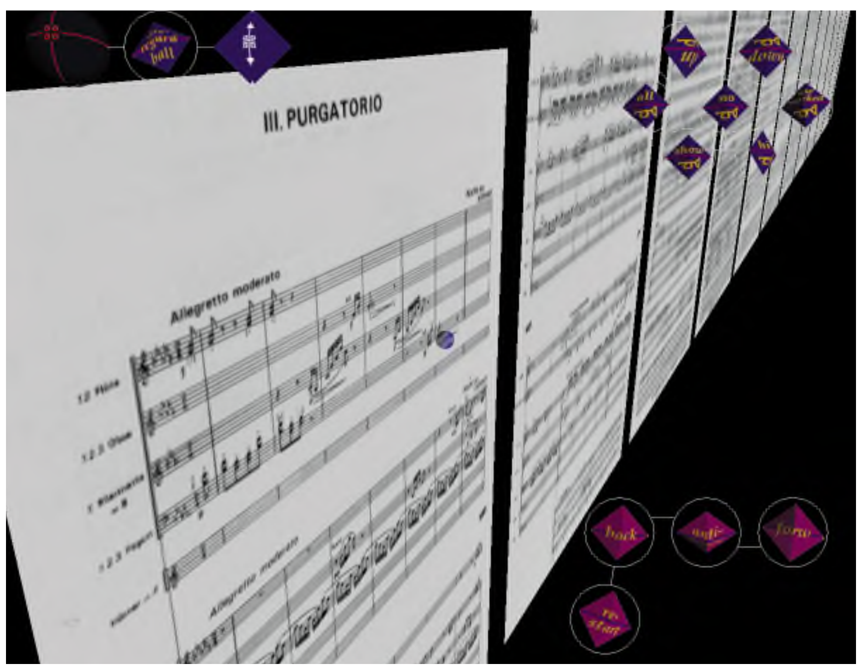

Contact
John Underkoffler
MIT Media Lab
Massachusetts Institute of Technology
20 Ames Street, Room E15-468B
Cambridge, Massachusetts 02139 USA
jh@media.mit.edu

CIMBLE: A Collaborative Learning Environment

The Franklin Institute

CIMBLE (CADETT Interactive Multi-User Business Learning Environment) enables six participants and a facilitator to enter a VRML world as 3D avatars and work together as a team on a guided task. The project is designed to integrate structured training methods into a sophisticated 3D world where the participants are no longer working in the same location. Since virtual, distributed work teams are replacing many traditional work teams, guidelines need to be developed to make this training effective in this new virtual environment. The CIMBLE prototype and project evaluation data highlight the results of a new effort in this area of online communities.

CIMBLE is a template for adult learning that accommodates a wide variety of training topics. The initial emphasis is on the soft skills required for effective teamwork. Other topics that may lend themselves well to this mode of learning include hazardous materials training, precision manufacturing, plant operations, military command and control, military field training (other than war scenarios), and group engineering and design applications.

The Consortium for Advanced Education and Training Technologies (CADETT) is funded by a DARPA grant. The consortium was formed to research advanced and emerging technologies, and study their use in workplace and workforce-development training programs. CIMBLE is one of three prototypes that the CADETT team is developing.

Contact
Edward Wagner
The Franklin Institute
222 North 20th Street
Philadelphia, Pennsylvania 19103 USA
+1.215.448.1119
ewagner@fi.edu
www.cadett.fi.edu

Collaborators
Sandy Testani
Edward Wagner
Sakunthala Gnanamgari
Judy Tumasz
Arne Jonsson
Karl Wehden
Dmitriy Skvirskiy
The Franklin Institute

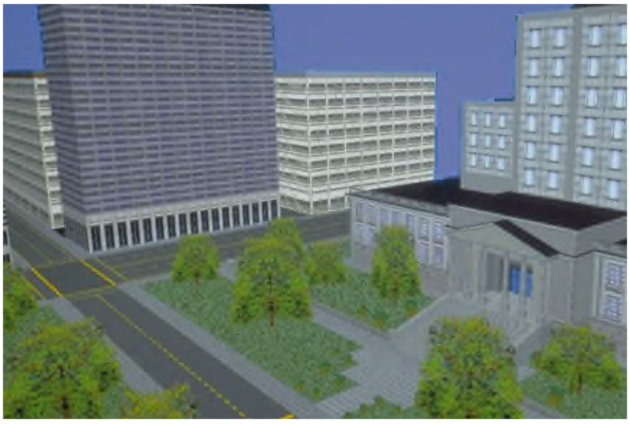

Isle of Write: Communication is the Landscape

Bell Laboratories

The Isle of Write is a dynamic VRML world for temporal messaging based on the metaphors of writing in the sand and skywriting. Adapted for the SIGGRAPH 98 community, it provides alternatives to the physical message boards, programs and calendars, hand-written posters announcing new Birds of a Feather (BOF) meetings, etc. typically used by attendees.

Over the last few years, we have been exploring new modes of interaction and developing metaphorical worlds whose primary landscape is comprised of communication objects. The Isle of Write is an island surrounded by an animated seascape. Its beach is the surface of several bulletin boards, and the sky is the surface on which broadcast messages are written. Visitors post messages using the SandTypewriter, and they dispatch the skywriting plane with the SkyWriter. Other conference materials populate the island, including paper abstracts printed on flying LiveWebStationery, BOF updates in SandCastles, and campfires fed by log statistics.

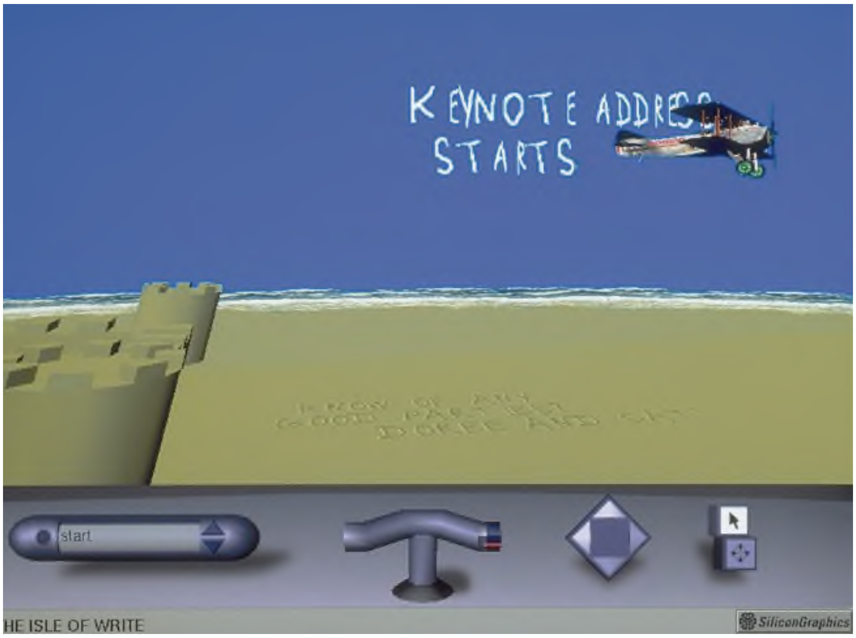

Contact
Dorèe Duncan Seligmann
and Cati Laporte
Bell Laboratories
Room 4G-608
101 Crawfords Corner Road
Holmdel, New Jersey 07733 USA
doree@bell-labs.com
cati@bell-labs.com

MicroDisplay Camera Phone

The MicroDisplay Camera Phone is an innovative, interactive demonstration created to illustrate the potential of the convergence of telecommunications, computer, and information technology.

At SIGGRAPH 98, attendees have the opportunity to make virtual calls with the Camera Phone, viewing real-time video on the phone's tiny display. When speaking normally into the phone, the user views live video images communicated to the phone's virtual viewer by the camera, which is directed to gather images from the user's field of view. Other visitors can participate in the virtual calls by viewing the same video feed on large, stationary monitors.

The MicroDisplay Camera Phone invites visitors to participate in a world where it is possible to transfer images, whether faxes or photographs, Web imagery or text, across distances not crossed by wires or cables.

Contact
Stephanie Silman
The MicroDisplay Corporation
3055 Research Drive
San Pablo, California 94806 USA
silman@microdisplay.com
www.microdisplay.com

Collaborators
Phillip Alvelda
Josie Ammer
Jeff Annis
Michael Bolotski
Steve Bristow
Dale Capewell
Mary Lou Jepsen
Eric Ko
Hongqin Shi
Rand Stadtman

MAGNET

France Telecom R&D

MAGNET, France Telecom's research and development project for streaming, interactive multimedia, is an implementation of VRML97 and MPEG4 for scalable platforms in telecommunications environments.

MAGNET will enable delivery of a media-rich environment over very low, consumer-available bandwidth such as 33K modems. Because it implements the VRML97 and MPEG4 standards, MAGNET represents a near-future technology that will be widely accessible to an Internet consumer audience, to business intranets and extranets, and to content creators. The compression capabilities in MPEG and binary encoding for VRML will demonstrate the exploitation of this low-bandwidth medium. The MAGNET architecture is scalable, and future work will include implementations for scaled-down clients such as laptops or smaller devices.

Our work presents an interactive multimedia scene to the user in an advertising context, such as one might find in the future Internet or long-distance markets. It allows both local and remote (server-side) updates of the scene and features 3D, 2D, audio, and video media running on Pentium class PCs with graphics acceleration. All media are encoded and streamed in 33K client bandwidth, including scene geometry, media content, and parameters for continuous animation.

Contact
Julien Signés
J. Jeffrey Close
France Telecom R&D
1000 Marina Boulevard, Suite 300
Brisbane, California 94005 USA
julien.signes@rd.francetelecom.com
jeffrey.close@rd.francetelecom.com

Collaborators
Michael Bourges-Sevenier
Renaud Cazoulat
Jeff Close
Yuval Fisher
Pascal Houlier
Olivier Ondet
Bick Thuy Pham
Didier Pillet
Julien Signés
Francois Ziserman

Jennifer James, Celebrity Auto Spokesperson

Jennifer James is a consumer-friendly, intelligent, interactive 3D ex-NASCAR driver who greets visitors at a virtual auto show and engages them in a dynamically customized five- to ten-minute dialog. Through natural social conversation, Ms. James elicits and offers information as she matches visitors to vehicles suited to their lifestyles and preferences. During this sales process, she applies dialog, facial expressions, and animated gestures to establish a relationship between her sponsor and each visitor. Ms. James exemplifies a new generation of interactive characters who will offer consumers a familiar and compelling sales experience.

Contact
Barbara Hayes-Roth
Extempo Systems, Inc.
650 Saratoga Avenue
San Jose, California 95129 USA
bhr@extempo.com

Collaborators
AnimaTek International, Inc.
www.animatek-int.com/
Extempo Systems, Inc.
www.extempo.com/

sigKIDS / Community Outreach

sigKIDS / Community Outreach

sigKIDS Art

Introduction

Kids: Reaching Out to Our Future

What would you give to see today's world through a child's eyes? Picture your sense of wonder, with all the digital tools for the imagination. Think how far your mind could soar, the worlds you could create, the places you could explore. In the last 25 years of discovery, computers have revolutionized the way we communicate and how we see the world. However, we've seen only a fraction of the full effect. Wait until we see what kids who have grown up with computers will accomplish in their lifetimes. They'll be free of the obstacles we faced in these early stages of technology. To see the future, we really need to look through the eyes of kids, and we need to help them show us what computer graphics and interactive techniques can really do.

This is the purpose of integrating the sigKIDS and Community Outreach initiatives, to reach out and affect the future of computer graphics through helping kids uncover its promise. We not only want to showcase the technological accomplishments of select kids, but also to inspire communities to offer all kids the opportunities this technology provides. In today's society, every child is at-risk. The magic of computer graphics offers the means to spark a lifetime of learning for even the hardest-to-reach students.

The exhibits and activities of sigKIDS and Community Outreach help show us how kids can empower themselves and overcome physical obstacles or melt social boundaries with innovative technology. Differently abled kids are obtaining better tools to communicate and interact with others. Disadvantaged kids can aspire to new career opportunities. This is possible when they have increased access to the technology plus guidance and inspiration from educators and professionals.

In addition to the conference activities, we worked within the local community to encourage collaboration among educators, professionals, and parents. We took the spirit of the SIGGRAPH conference beyond the walls of the convention, beyond the pages of proceedings, beyond the one-week event.

Community Outreach: Bridging Education and Industry

We invited educators and professionals to collaborate in creating activities that explore the potential of new technology for kids. We found a unique role for SIGGRAPH 98: providing a crucial bridge for educators and industry professionals to come together and better prepare young students for a technologically intense future.

The digital revolution is changing the face of education almost as much as it is changing industry, and it is extremely difficult for teachers to keep up with the changes. The Community Outreach Team discovered a strong need that SIGGRAPH 98 could fill within the community. Teacher organizations came to us to give their members a deeper understanding of computer graphics as an educational tool, to offer insight on the ever-changing careers in the industry, and to share with students the relationship between an exciting career in computer graphics and their everyday studies of math, science, language, and art.

We developed partnerships with these groups. We gave presentations, facilitated roundtables, organized hands-on labs, and provided mentoring sessions for a variety of state-wide organizations. These associations included the Florida Art Educator Association (FAEA), School-to-Work (STW) Conference, Arts for a Complete Education (ACE), Florida Film Institute for Education (FIFE), Boy Scouts of America, Florida Association for Media Educators (FAME), and the Florida Assistive Technology Association (FATIC/FDLRS), to name a few.

Insight into computer graphics industries helped teachers overcome one of their biggest challenges – motivating students to learn. They found that they are able to highlight exciting careers that students could imagine themselves following. When they meet successful professionals in their own communities, kids gain new motivation to succeed in school.

We partnered with schools and organizations within the local community to create programs such as the sigKIDS Outpost to offer a hands-on outlet for a diversity of participants (from valedictorians to at-risk kids). The teams worked side-by-side with professionals. Local and international companies contributed mentors and donations. Once a month, parents were invited to support their kids' journeys with computer graphics. The motivated and self-directed students brought these workshops to educational conferences to encourage other teachers throughout Florida to start programs.

This connection between industry careers and education sparked our relationship with the Florida Department of Education's (FDOE) School-to-Work program. Together, we found a way to send representatives from every county in Florida to see for themselves the excitement of the computer graphics industry and the spirit of SIGGRAPH. With a matching grant from FDOE and the SIGGRAPH Education Committee, a four-member team was selected from 67 counties to receive mini-grants to SIGGRAPH 98. Each team includes a cross-section of the education community, including a teacher, an administrator, a parent, and a student.

To encourage participation from the industry, we also partnered with professional and civic organizations in special events. These included the Art + Technology festival with the Orlando Museum of Art, Nickelodeon Studios, ResFest Digital Video Festival at the Enzian Film Organization, the Central Florida Computer Graphics Association Conference, and the Winter Park Art Festival.

Community Outreach 98 and Beyond

"Linking grand goals to realistic scenarios for accomplishing them takes impassioned imagination combined with scientific rigor. There are so many important problems to work out, there is room for everyone to contribute."

- Ben Shneiderman
Between Hope and Fear
ACM's Beyond Calculation

Technology has made our lives more exciting, our careers more expansive, and the world more accessible. Technology can also bring a new level of confusion, anxiety, and alienation to those who do not have access to or understanding of the new innovations. How can we as a society of artists, scientists, educators, and visionaries reach out to communities, share the spirit of SIGGRAPH, and turn the overwhelming demands of the future into visions of hope? We have only started finding answers; it is our hope that this effort continues to grow stronger in SIGGRAPH's future.

sigKIDS/Community Outreach Committee

Co-Chair
Adele Newton
Newton Associates

Co-Chair
Chris Stapleton
Universal Studios

Co-Director
Scott Lang
The Academy for the
Advancement of Science and Technology

Co-Director
Anne Richardson
O.C., Inc.

sigKIDS Animation Festival Coordinator
Kevin McTiernan
The Academy for the
Advancement of Science and Technology

Heidi Dunphy
Independent Animator

Nancy Krebsbach
Evans High School

Mitchel Groter
Video Central

Steve Schain
O.C., Inc.

Lynn Finch
Finch Interactive

Chris Carey
Orange County Public Schools

Darlene Wolfe
River Ridge High School

Pat Johnson
The Art Institute of Fort Lauderdale

Sally Brahier
Former District Liaison, PTSA

Rob Brahier
River Ridge High School

sigKIDS Jury

Heidi Dunphy
Independent Animator

Tim Comolli
South Burlington High School Imaging Lab

Maria Roussos
University of Illinois at Chicago

Chris Carey
Orange County Schools

Alice

Jeff Pierce
Tina Cobb
Randy Pausch
Carnegie Mellon University

Alice is a development environment for the creation of interactive 3D worlds. Our primary goal is to make the program easy to learn and use for non-engineering junior high, high school, and undergraduate students. We want students with little or no programming experience to be able to run through a 30-minute tutorial and start building fun and interesting 3D worlds right away. Alice is available free for Windows 95 and NT platforms. The latest version is available at:

alice.cs.cmu.edu

The Alice project was started by the User Interface Group at the University of Virginia, which later moved to Carnegie Mellon University and became the Stage 3 research group. The program is an interactive 3D graphics authoring system with a straightforward scripting environment, a powerful animation facility, and an intuitive overall design solidified by extensive user testing. The design of Alice's user interface was heavily influenced by cognitive and perceptual psychology literature.

Alice is currently in widespread use. It is used in the Lynchburgh, Virginia school system and has been used in class projects at the University of Virginia and the University of Central Florida. Parents are also using Alice to introduce their children to 3D graphics. Our hope is that, just as LOGO introduced students to 2D graphics more than a decade ago, Alice will introduce a new generation of students to building interactive 3D worlds today.

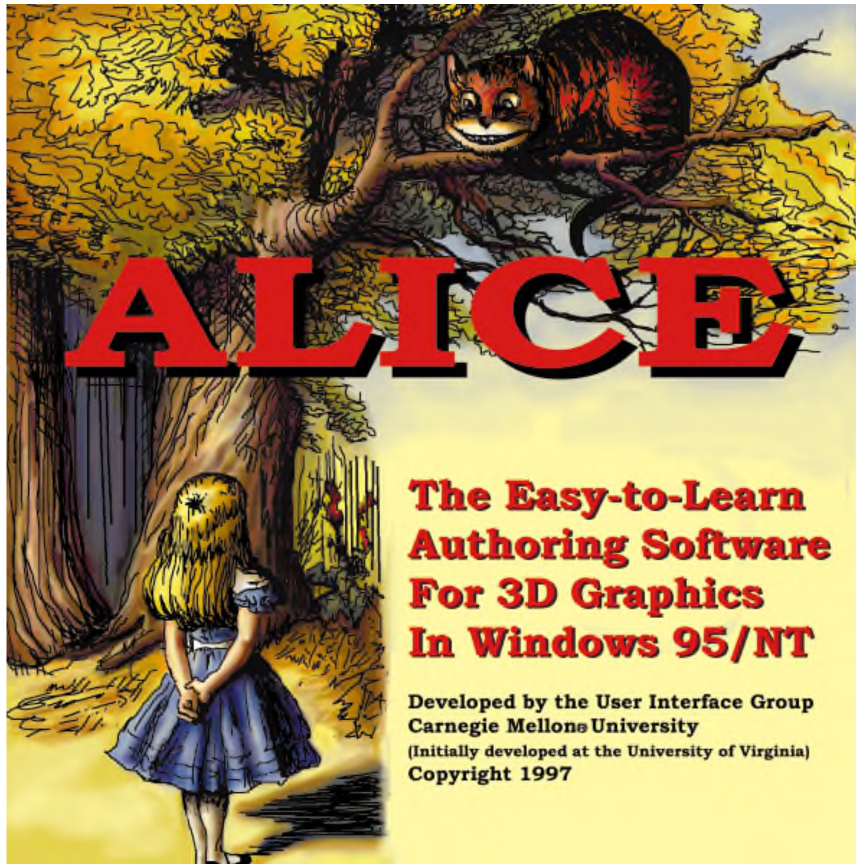

The alpha version of Alice (over 6,500 copies distributed) was released at SIGGRAPH 96. The beta version (over 20,000 copies distributed) was released at SIGGRAPH 97. The program was used to create virtual experiences in the Digital Bayou at SIGGRAPH 96 and the Electric Garden at SIGGRAPH 97.

Contact
Tina Cobb
Computer Science Department
Carnegie-Mellon University
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213 USA
+1.412.268.8724
tinac@cs.cmu.edu

Binary Biker Project: An Exploration of Motorcycles, Art, and Technology

Rick Barry
Pratt Institute

In support of the New York City ACM SIGGRAPH Dan Preda Scholarship Fund

In July of 1998, a motorcycle convoy led by two specially rigged motorcycles and a support vehicle departed from the campuses of Pratt Institute and the School of Visual Arts in New York City, to begin an event-filled journey to SIGGRAPH 98 in Orlando. These Binary Bikers were accompanied by an invited entourage from the worlds of art, technology, education, and motorcycling.

Contact
Rick Barry
Pratt Institute
Department of Computer Graphics
200 Willoughby Avenue, ARC, LL, F-10
Brooklyn, New York 11205 USA
+1.718.636.3411
+1.718.399.4494 fax
rbarry@interport.net

The lead vehicles and Binary Bikers were equipped with live audio and video set up to transmit two-way, time-delayed sound and pictures via satellite or cellular connection to a central Web server at Pratt Institute. This Web site (binary-biker.org) was the focal point for the Binary Biker Project. Remote participants were able to log in and "ride along" as "virtual bikers."

During the journey, still and moving images were captured by both the Binary Bikers and the virtual bikers via the Web site and digitally stored on the Web server in an online "image pool" for common access by all participants. Artists were invited to capture, download, and access these images, use them to create their own original imagery, and upload them to an online "submission pool" for subsequent judging. The winning selections were placed in an online Binary Biker gallery, and a limited number were selected and uploaded to the SIGGRAPH 98 Guerrilla Gallery, where they were output to color imaging devices, framed, and exhibited in a Guerrilla Gallery Binary Biker showcase.

At SIGGRAPH 98, the Binary Biker exhibition area contains several workstations offering access to the Binary Biker Web site, as well as computer graphics software and hardware tools to enable artists to create original works and upload them to the SIGGRAPH 98 Guerrilla Gallery.

binary-biker.org
An Exploration of Motorcycles, Art & Technology

CAROL
(Culture and Arts of Rochester Online)

Gordon Goodman
Stephen Jacobs
Rochester Institute of Technology

Four years ago, students in “Topics in Interactive Multimedia” built a Web site for a Rochester, New York “living history museum.” From that initial site, CAROL has grown to include over 20 local arts and cultural institutions’ Web sites, an online events calendar, and a database of local artists. CAROL is also a consortium of organizations actively involved in shaping the future and growth of the sites. CAROL’s model – bringing non-profit cultural organizations, industry, and a university together – has been recognized by both the ACM (which invited CAROL to exhibit at ACM ‘97) and the NEA (which granted the local arts council $30,000 to support the project).

The course brings students together with representatives of arts organizations to develop or upgrade Web sites. The representatives attend the class as both clients who support site development and as students learning to maintain their sites. The course continues to add new sites and update old ones, breaking ground in electronic commerce and providing a source of undergraduate independent study and graduate thesis projects.

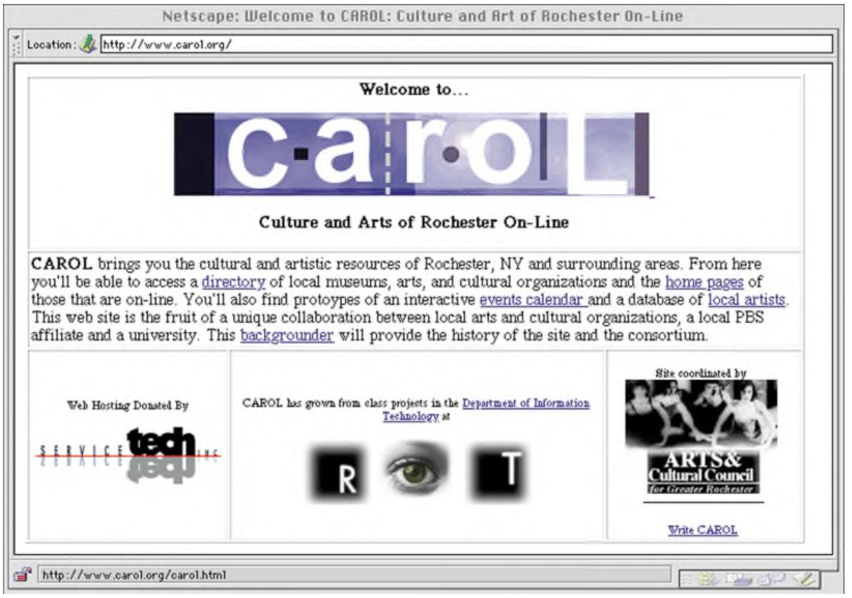

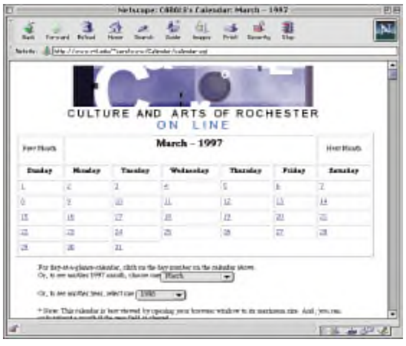

CAROL goals:

1 To create a common point of entry and an interactive cultural events calendar on the Internet for Rochester’s cultural and arts organizations.

2 To provide Web space and expertise for groups that want to establish a Web presence but do not have the resources to create one on their own.

3 To provide a long-term scheduling resource that allows organizations to see a common calendar of events across institutions.

4 To act as a clearinghouse, connecting members of industry and volunteers who want to support CAROL members with the appropriate contacts.

5 To investigate the possibilities of an organization-wide intranet, allowing such groups access to networked technology.

6 To provide students with interesting, exciting, real-world projects as part of their academic experiences.

Contact
Stephen Jacobs
Rochester Institute of Technology
Department of Information Technology
102 Lomb Memorial Drive
Rochester, New York 14623 USA

The Cyberarium

Dave Warner
Institute for Interventional Informatics

The Cyberarium is an innovative, integrative environment that stimulates discovery and exploration of creative and socially enriching interactive techniques through intelligence-engaging, imagination-inducing, hands-on experiences with novel human-computer interaction devices.

The Cyberarium has three major areas:

1. A Widgiteering Lab, where kids get hands-on exposure to and experience with hardware and software tools that will enable them to build and test creative and novel interactive tools. Through demonstration and discussion, they are also exposed to the interface problems of the disabled computer user. The software component features "Neattools," an object-oriented, multi-platform, free program designed to take data from a wide variety of input devices and use it to control any output (music, joystick and mouse output, network connections, and remote control vehicles, for example).

2. A space where attendees can experience and learn about new communication technologies and interactive techniques designed to improve the quality of life, this "classroom" educates kids in the Widgiteering Lab about socially relevant applications. It is also a hands-on demonstration area that showcases accessible interfaces for disabled children and demonstrates how the Internet can be used for telemedicine, to promote good health and recreational activities that can reduce stress, elevate mood, and maintain wellness

3. A real-world/hands-on experience for educators and parents. In parallel with the Widgiteering Lab, The Cyberarium provides workshops and seminars designed to empower teachers and parents so they can

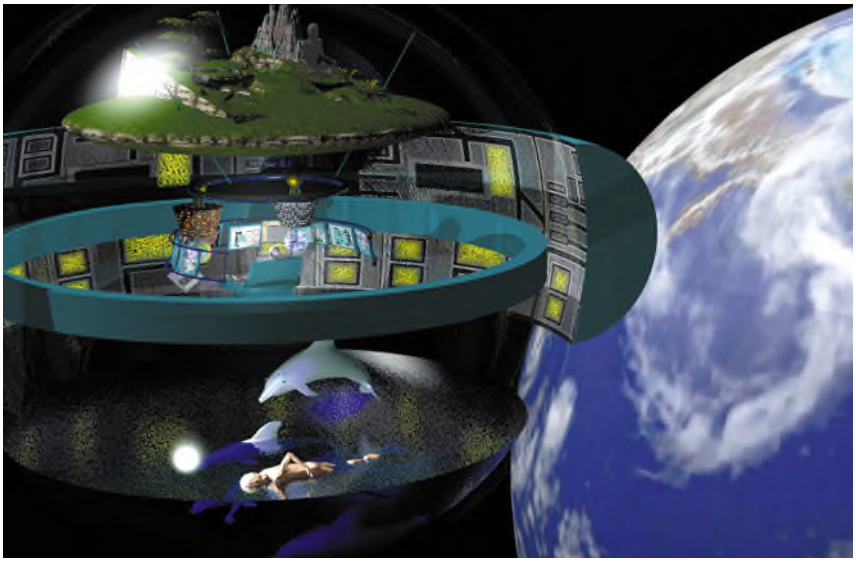

easily utilize emerging communication tools and interactive technologies in the classroom and the home. Emphasis is on methods and tools for teaching the science, math, and social skills that will be needed by the "citizens of cyberculture" they are educating and raising.

The Cyberarium's intent is to develop an experimental environment in which "cultural rapid prototyping" can be observed, researched, and refined – a place where social dimensions can be intelligently combined with the rapid development of the information infrastructure.

Contact
Dave Warner
Institute for Interventional Informatics
500 University Place
Syracuse, New York 13210 USA
+1.315.423.4676
+1.315.443.1973 fax
davew@npac.syr.edu

Dream Map Tapestries

Nancy Krebsbach
Mitchel Groter
sigKIDS Outpost

This event celebrates the work that sigKIDS groups everywhere have done throughout the year. The kids' work is presented on videotape. After the screening, a panel discussion takes place among people from industry and education who truly care about helping kids and their parents. Topics addressed by the panel include:

- Creativity
- How I made it in the business
- Portfolio development
- Requirements to get into and secure scholarships at leading schools
- Internships, externships
- The joy of creating a professional life that you truly love

After the panel discussion, attendees have access to a career fair, school expo, portfolio reviews, and special events that focus on the anatomy of an animation project and the history of three artists' journeys into their profession.

The Dream Map Tapestries connects motivated sigKIDS with mentors who will continue the relationship beyond SIGGRAPH 98. The kids have a unique opportunity to talk one-on-one with people who can be role models for them. They hear, directly from the professionals, what it took to become successful in the computer graphics industry, what particular talents and skills they needed, what personal traits are the most helpful, and exactly how young people should prepare themselves for a successful career in computer graphics.

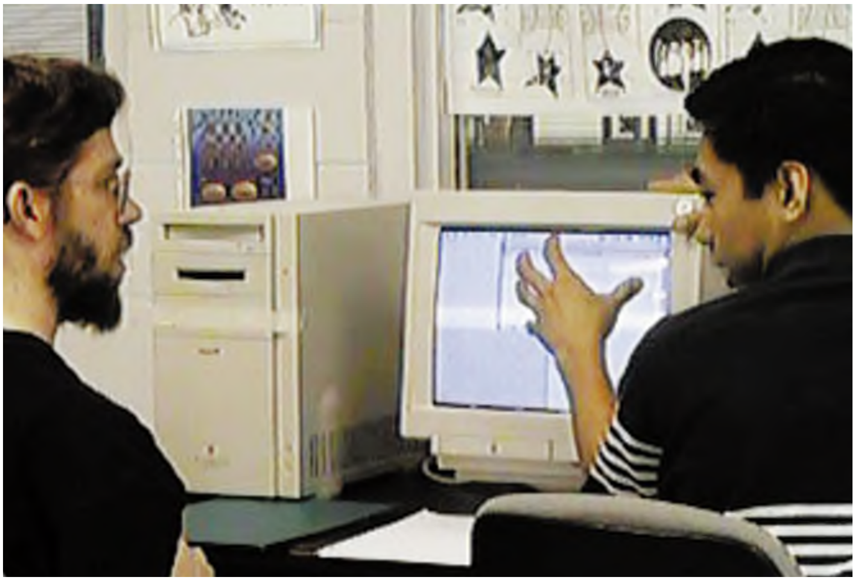

Contact
Mitchel Groter
Video Central
825 Sweetwater Island Circles
Longwood, Florida 32779 USA
+1.407.786.0017
Pathways@parkave.net

Nancy Krebsbach
Evans High School
858 Park Lake Court
Orlando, Florida 32803 USA
+1.407.293.4900
Faea@magicnet.net

Express Link-Up

Mark Jones
Pat Ryan
DeMonfort University

Founded in January 1997, Express Link-Up empowers hospitalized children by providing appropriate computer and communications technology, including a dedicated secure intranet that enables them to play, learn, communicate, and develop critical life skills. The NHSnet, which supports the intranet, is the communication network of Britain's National Health Service. Software that spans education (including virtual reality interactive chemistry, encyclopedias, literacy software, and word processing), email, Web-site building, and games have made the system an enormous hit with patients and hospital school teachers. Patients can send and receive email from their schools and friends, but a high-level firewall blocks unwelcome contact. Speech-activated software assists patients who have problems using a keyboard. This is of great value for patients with injured hands, those suffering from cerebral palsy or head injuries, or even young children who have not mastered the keyboard. 250,000 children pass through UK hospitals each year, of which 150,000 are considered long-stay. Some children (for example, those who require six-hour kidney dialysis three days a week) suffer terrible disruption to their education. With PCs next to their beds, they can bring work to the hospital, do school projects, study for exams, or use relaxation software to overcome stress.

Contact
Pat Ryan
Express Link-Up
32 Matham Road
East Molesey
Surrey KT8 0SU, United Kingdom
+44.181.941.0102
101521.1516@compuserve.com

Kerry using the PC in the ward, watched by her sister and a director of BT Syntegra (a contributor to the system).

Tasharn is hooked up to his chemotherapy drip and using the PC. Alongside him is Chris Jarvis (a BBC Children's TV Presenter), watched by Mandy Boydle (Teacher in Charge of the Hospital School) and Ian Taylor (Member of Parliament – the Minister who helped the project get started).

St. Bartholomew's Hospital in London, where these photographs were taken, has just celebrated its 875th birthday. There has been a hospital on the site since 1123, which makes it the oldest hospital in Britain.

Generation LEGO

Christian Greuel
Lisbeth Frølunde
LEGO

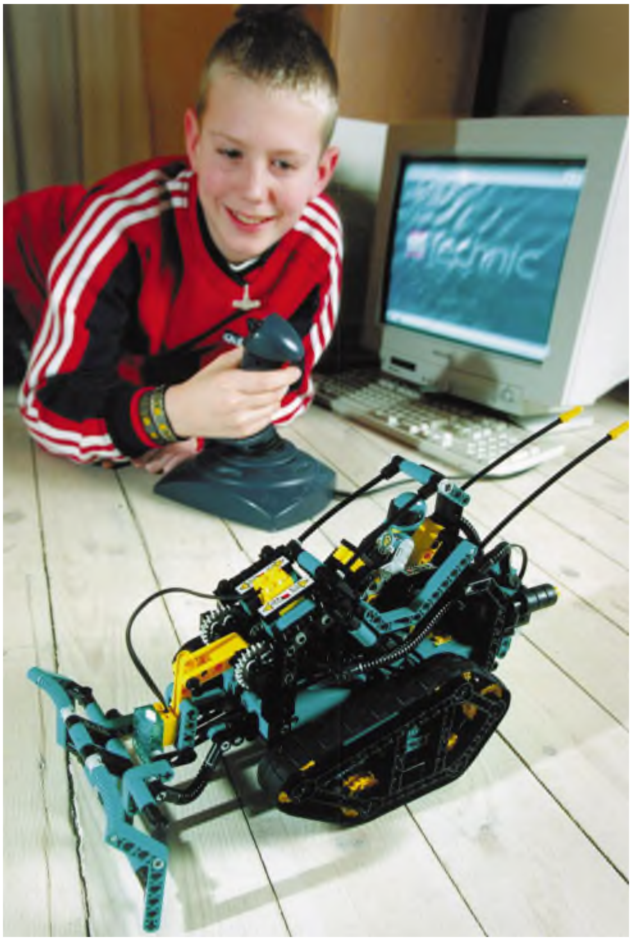

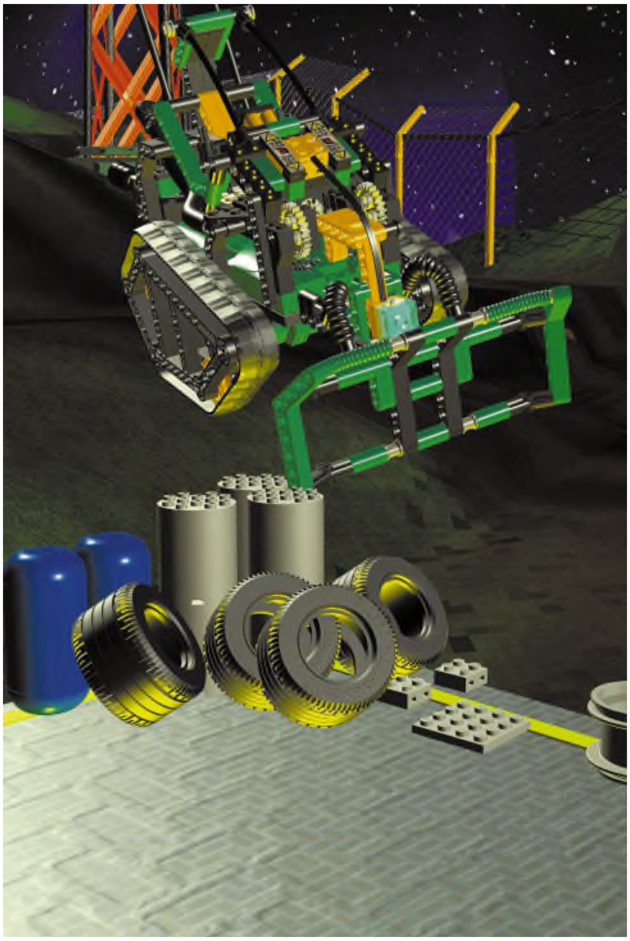

Just imagine a child's toy jumping off the shelf and taking on a life of its own. Through the magical combination of toys and technology, this dream has finally become a reality. The LEGO Group has developed a new generation of intelligent construction toys and learning tools for children. The key element is a programmable LEGO brick equipped with a microchip that enables children 10 years old and up to give behaviors and personalities to their hand-built models.

Two new LEGO products (LEGO Mindstorms Robotics Invention System and LEGO Technic CyberMaster) bring construction toys into the information age and give children an opportunity to build their own intelligent and interactive inventions. LEGO Mindstorms, developed in cooperation with the MIT Media Lab, is the first in a new generation of challenging, creative learning tools that enable children to use a PC to program intelligence into their LEGO inventions. LEGO Technic CyberMaster is a futuristic play set that combines the virtual fun of on-screen adventure with physical models. Children can build their own models and bring them to life with a home computer.

These new LEGO products are designed to empower children in an age of increasing computerization and challenge their creativity, craftsmanship, programming, and critical thinking.

Contact
Lisbeth Frølunde
Concept Developer
LEGO A/S, SPU-Darwin
Klovermarken 120
DK-7190 Billund, Denmark
+45.75.33.11.88 x 5172
+45.75.35.47.19 fax
lisbeth@digi.lego.com

HistoryCity

Terry Lim
Kent Ridge Development Labs

This 3D virtual world for kids is set in 1870s Singapore. The world features functional buildings; personal decorate-able rooms; animated objects that can be picked up and dropped; personal theater stages upon which dioramas can be built from objects found in the world; over 200 avatars representing people and occupations of the period; agents that provide news, stories, poems, jokes, pawn brokering services, and messaging services; clubhouse memberships; maps; and 22 communities, each with its own distinct architecture, and music.

The world has been tested in both Singapore and Canada to study users' individual and social behavior and learn how to make HistoryCity more engaging. Kids can also contribute material to the community's ongoing corpus of news, stories, poems, and jokes, including audio contributions. Multilingual text chat, as well as audio chat (using microphones and sound cards), are supported.

Contact
Terry Lim
Kent Ridge Development Labs
21 Heng Mui Keng Terrace
Singapore 119613
terry@krdl.org.sg

Hyperscratch ver.7

Haruo Ishii

This interactive piece allows users to generate a variety of sounds and images as if they were using simple hand motions to paint a picture on a canvas of space or playing a piano with invisible keys. Instead of providing an interactive display through which users move in a prearranged, simulated environment, this piece offers a way to spend time more creatively in a physical space.

No visual interface or control device is used to operate Hyperscratch. Users simply stand before the screen and move their hands at face level to generate and change images and sounds. In other words, an invisible three-dimensional interface allows manipulation of images and sounds. Infrared light is cast upon the moving hands, and highly sensitive video cameras capture the hand motion. The computer processes the motion to generate images, and a MIDI signal is sent to the digital sound sampler to generate sounds.

This piece is designed to provide users with uninhibited creative space and time. The only input devices are hands and bodies. Both hands can be used to create sounds and images, which allows users to move as if they are conducting a symphony or dancing. Such natural and free body movement is not possible through a mouse, a touch panel, or a space-input device. Thus Hyperscratch is not simply an input or pointing device but a unique environmental interface that mirrors the physical motion of the user.

Those who expect to experience some preconceived notion of virtual reality or an interactive video game will be disappointed. This piece affords users the freedom to create their own music and art on a large canvas, just like the abstract expressionist painters.

Contact
Haruo Ishii
30-1 Ishihata Naruomi-cho Midori-ku
Nagoya-shi, Japan 458-0801
+81.52.622.8697 voice/fax
MXC00275@niftyserve.or.jp

InterSpace Station

Robert Rothfarb
NTT America Inc.

InterSpace Station is a networked 3D virtual world environment that provides a shared laboratory. Students are able to conduct science experiments, interact with each other, and participate in unique presentations about astronomy, physics, and life aboard an orbiting space station. The InterSpace 3D multi-user virtual world client/server software platform from NTT allows users, as avatars, to navigate digital spaces and communicate with each other using real-time audio, facial image video, and text. Students and teachers participate from SIGGRAPH 98 and two remote sites, one in the San Francisco Bay area and one in Japan. Visitors may also participate via the Internet. NTT plans to collaborate with educators and SIGGRAPH Professional Chapters to allow students to participate in this unique interactive distance-learning environment.

Contact
Robert Rothfarb
Virtual World Designer
NTT America Inc.
100 Shoreline Highway, Suite 100A
Mill Valley, California 94941 USA
+1.415.289.7729
+1.415.289.4621 fax
rob@nttlabs.com

KidCast For Peace: Solutions For A Better World

Peter H. Rosen
Visionary Artists Resources Including
Other Unique Services

Children of all ages share their art live and direct us to their KidCast For Peace Web sites, VRML worlds, chat rooms, and interactive music spaces. Children at each participating site respond to comments and questions from local and cyber audiences. Free CU-SeeMe videoconference software enables real-time visual and audio interaction. The video camera focuses on the child and art together, then zooms into the art. A KidCast Central (creativity.net/kidcast2.html) moderator encourages the other sites to respond to what they see, evoking questions, feelings, and impressions. A moderator at each site takes responsibility for the following:

- Identifying the site. For example: "This is John in Honolulu. We have a question from our audience."

- Having a microphone available to provide a sound feed, (or type that scrolls across the CU-SeeMe window – visible to all active participants, and "lurkers" who just watch and listen).

- Introducing the artists and asking them to explain their work. These questions will catalyze KidCast community interaction.

The world needs an infusion of new ideas and loving connections to defuse the downward spiral of destructive human interactions. KidCast For Peace is part of a developing network of physical Creativity Cafes, other "New Schools," and evolutionary organizations that are drawn toward cooperative activities. This live and interactive global community forum enables all of us to suggest how we might heal our differences, hear our needs, and feed our spirits as a united humanity. KidCast taps into the primal creative force so we can listen to possible

solutions to personal and planetary problems offered by kids from their more "untainted" point of view.

The vision of KidCast for Peace is: to connect children in schools and homes around the world and facilitate their ability to contribute to the world they will inherit; to inspire educators and parents to guide children (of all ages) in expressing their feelings for how to make this a happier, healthier, safer and more peaceful world; to create a global network of kids in physical locations that connect in regularly scheduled live and cyberspace broadcasts on the Web; and to build a KidCast For Peace Web site and "Peace Place," a repository for solutions that contribute to a better world using the artistic expressions created by kids.

Contact
Peter H. Rosen
Visionary Artists Resources Including
Other Unique Services
2263 Sacramento Street, #2
San Francisco, California 94115 USA
+1.808.573.3943
peter@creativity.net

Living and Learning

Candis Hoffman-Bomse
Erin Hethington Ebert
Savannah College of Art and Design

This project provides instruction on a critical health issue affecting today's youth: AIDS. It uses art to educate teens about risky behaviors, inaccurate myths, and subject matter that is often considered taboo in person-to-person discussions.

As artists, we recognize art as a form of communication and wish to explore our understanding of the various modes of digital expression to further human knowledge, or at the very least to present new ideas to others. One of the ways we are attempting to do this is to develop new and engaging ways of conveying important information to groups of people who may be difficult to reach or resistant to the information when delivered in conventional forms.

In Living and Learning, we research cultural issues, media conventions, and technological resources in the hope that this postmodern approach will help convey much needed information about AIDS to an age group that sees itself as unaffected and invincible. Our goal is to engage this group long enough for them to become involved in the interplay of moving imagery and sound, virtual games, and navigation of the work's topography via mouse-driven interactivity. The result, we hope, will be the work's ultimate goal: education, the best defense against AIDS.

Contact
Candis Hoffman-Bomse
Erin Hethington Ebert
Savannah College of Art and Design
210 E. 60th Street
Savannah, Georgia 31405 USA
cbomse@worldnet.att.net
E2ebert@aol.com

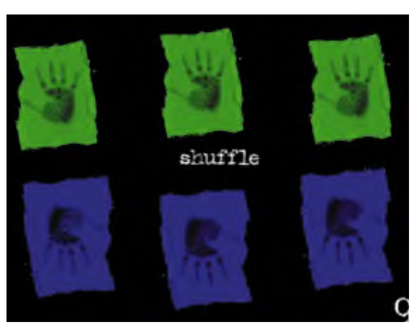

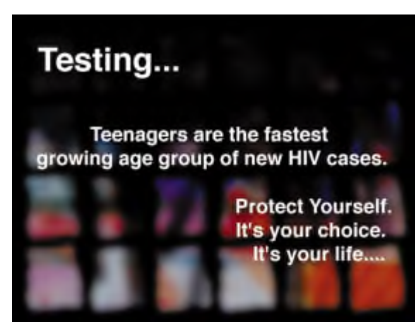

Mario Net

Celine Jaspart
Animaçáo-IRIT

Mario Net is a show with virtual puppets animated in real time by puppeteers using Polhemus sensors. All the animation sequences are video-projected on a large screen or adapted screen. In this new kind of spectacle, the virtual characters improvise with the children in the audience like real puppets or actors, through real-time animation systems.

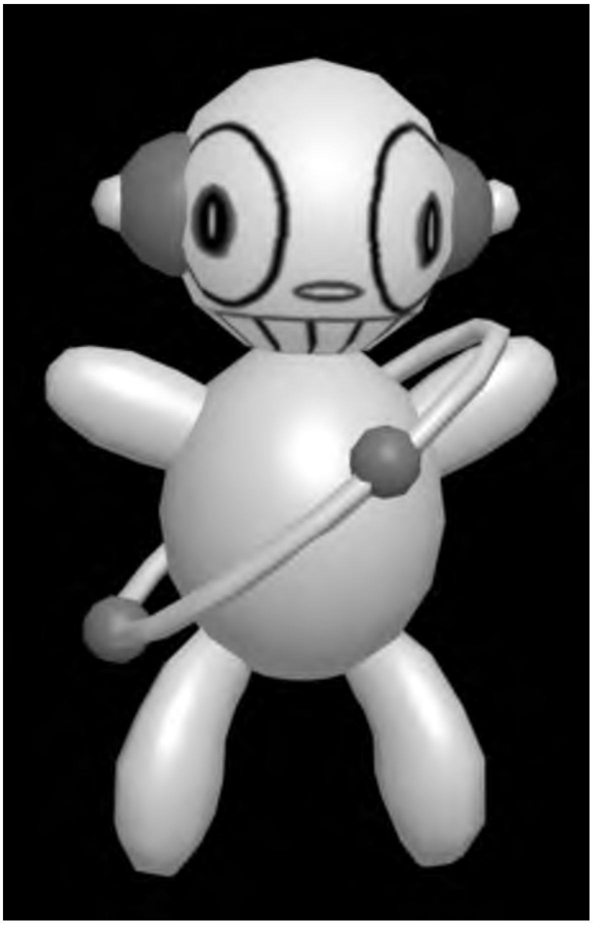

Our interest is to animate our virtual puppets like real ones, to allow them to find their own rhythms and personalities, and to create a new sceneography around computer technologies, a new space. Mario Net mixes many different personalities: actors, puppeteers, computer artists, and computer engineers. In this collaboration, artists work with technological constraints, and computer engineers discover strange applications of their work.

Contact
Celine Jaspart
Animaçáo-IRIT
14 rue Saint Bertrand
31 500 Toulouse, France
+33.5.61.20.21.26
jaspart@worldnet.fr

The Virtual Archaeologist

Nobuo Masuda
Cyber Entertainment, Inc.

Developed for an archaeological museum in Nara prefecture, Japan, this kiosk program introduces young visitors to the world of archaeology. Nara prefecture is well-known for its historical artifacts. In ancient times, it was the capital of Japan. One of the most important objectives of this software is to inspire newer generations and help them understand the value and significance of archaeological studies.

At this type of museum, visitors usually see already-unearthed artifacts. The real excitement in archaeology, however, is in the actual process of searching out historical clues. In this interactive kiosk, visitors experience this “excitement” as they become detectives to solve “historical mysteries buried in the ground.”

The kiosk uses “push” technologies in an “interactive document.” Sometimes the software automatically opens up new pages and spontaneously conveys information to the users, who then feel as though they are discovering clues and mysteries.

The software includes eight historical periods from the primitive to the modern age, though it emphasizes the Age of the Ancient Burial Mounds (third to sixth century AD), when Nara was the center of Japan.

Contact
Nobuo Masuda
Cyber Entertainment, Inc.
5111 Denny Avenue #10
North Hollywood, California 91601 USA
+1.818.505.1837
+1.818.505.1548
masuda@cyber-net.co.jp

The Virtual Art Gallery/ Streaming Video on the Web

Ann Ioannides
Colette Stemple
Coral Reef Senior High School

Imaginative students from Coral Reef Senior High School (Miami) created two projects:

- A VRML art gallery, featuring works of art produced entirely by first- and second-year students.

- A streaming video newscast. This "Cudavision," five-minute daily newscast has been digitized for worldwide viewing. Tune in frequently to what is happening at the school!

Thanks to state-of-the-art equipment, the guidance of teachers Ann Ioannides (Television Production/Journalism) and Colette Stemple (Photography), and Victor J. Deleon of Digitalo Design and Florida Atlantic University, the students created these unique projects. The school features six magnet programs as well as a full athletic and extra-curricular program, attracting the best and brightest students from Dade County. The 1997-98 year is the school's first year of operation.

crhs.dade.k12.fl.us

Contact
Ann Ioannides
Coral Reef Senior High School
10101 SW 152 Street
Miami, Florida 33157 USA
+1.305.232.2044 x 322
+1.305.256.1697 fax
ioannia@mail.firn.edu

When Children Draw In 3D

Katiuska Varela

In this project, children's drawings are used to create a 3D animation. It begins with a workshop that asks children a series of questions about a story: Who is this story about? What is it about? Where does it develop? The workshop results are used to prepare a script, then the script is adapted for a storyboard, and the children draw the characters and places where the story happens.

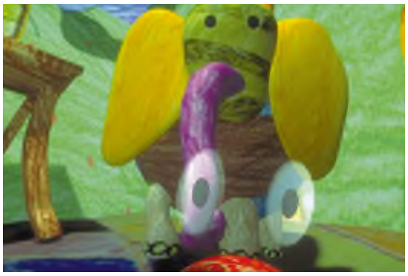

Portions of the children's drawings (for example, the legs of one drawing, the head of another, and the ears of a third) are combined and scanned, then colored by the children. The colors are used as textures for elements of the story, and the scans are used to make the 3D forms.

Finally, the storyboard scenes are prepared for animation, and characters, locations, cameras, lights, and other elements are organized to form the final edition of the story.

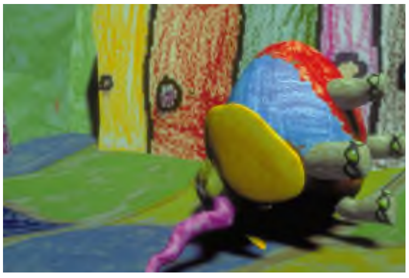

Contact
Katiuska Varela
4 rue Calmels
75018 Paris, France
+33.1.42233862
+33.58.28611646 fax
gonzalez@ensba.fr

Where Stories Meet by TeleCommunity

Robert Dunn
TeleCommunity Project/College of Liberal Arts
Duquesne University

Where Stories Meet by TeleCommunity represents a confluence of ideas, images, and experiences expressed through digital media and shared over the Internet by students from Jerusalem, Istanbul, Orlando, Pittsburgh, New Orleans, and Los Angeles.

In studio environments, young people (11-17 years of age) work on computer graphics, animations, Web sites, and multimedia projects that evolve into narratives about cultural and social issues, personal and family histories, fiction and fantasy, and documentary efforts. Stories, shared among individuals beyond boundaries, test the perception of reality and lend impetus to the creation of new myths and visions.

Underlying the social fabric of this project are relationships enhanced by personal meetings, dialogue, and negotiated understandings between participants. This is sustained remotely by periodic and ongoing network contact, email, Internet videoconferencing, Web site updates, and file exchanges. Over-arching themes bring diverse responses and help to spark the imagination. The individual grows and experiments within the collaborative group setting, and energies are shared with remote partners.

The Pittsburgh node of TeleCommunity invites young people to a studio setting at the Duquesne University College of Liberal Arts. The project is extended to an international level by initiating dialogue and building working partnerships with interested educators and artists in other countries.

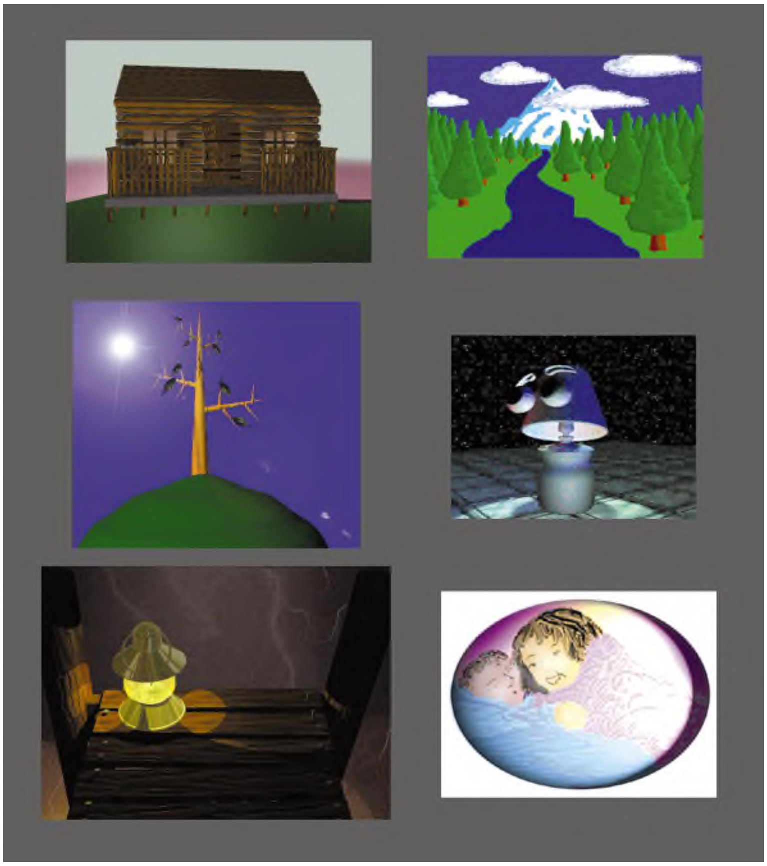

Teachers, artists, art educators, and consultants volunteer their time to contribute to the undertaking. An atmosphere of conceptual exploration and collaboration among peers nurtures the creative process. The Web site (www.telecommunity.org) reflects works in progress, resources, and past efforts.

Contact
Robert Dunn
TeleCommunity Project/College of Liberal Arts
Duquesne University
111 The Oaks
Pittsburgh, Pennsylvania 15215 USA
+1.412.781.1563
+1.412.781.8138 fax

Save the Planet: Eco-Art on the Web

June Julian
New York University

Do you want to save the planet? The Web is the way! Students from around the world are participating in A World Community of Old Trees (www.nyu.edu/projects/julian/). This project is a good working example of how the World Wide Web can be used as an interactive medium for communication on global ecology issues.

Since this project was first launched two years ago, students from around the world have been sending in their art work, poems, stories, and eco-facts about the special trees in their environment. So far, 21 U.S. states and 18 foreign countries are represented on at least 100 separate Web pages. The youngest contributors are Illinois kindergartners, and the oldest students are in the twelfth grade in Nevada. One New Jersey boy even designed his own interactive space on the project, where he asks the world to manipulate his pictures and send them back to the site. In fact, the project is especially designed to welcome your own unique ideas about the old trees of the world, so send them in! We still have lots of states and countries that are not represented yet. What is the oldest tree in your neighborhood?

Contact
June Julian
New York University
P.O. Box 81
Gladstone, New Jersey 07934 USA
+1.908.234.1611 x 231
julianj@acf2.nyu.edu

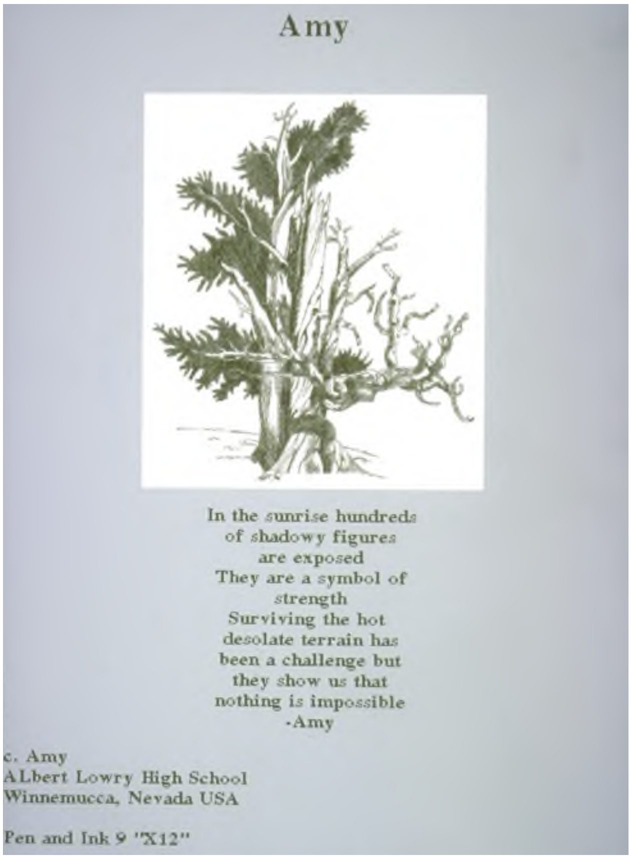

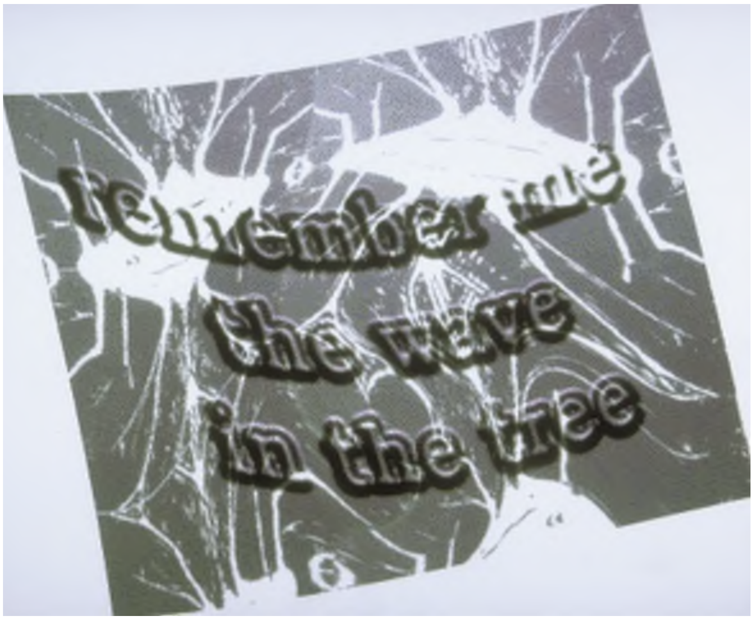

The sigKIDS Art Show features 30 pieces selected from over 150 submissions. The students who created these images come from a variety of backgrounds. Some live in rural communities, others in large metropolitan areas. Some attend public schools, and others attend private schools. Some are in elementary school, while others are one step away from college.

The common thread that ties these diverse groups together is the computer and its use as a tool in the creation of beautiful imagery. Much as pencils and charcoal can be found in any art student's toolbox, access to a computer for artistic expression is becoming more commonplace for younger students.

Both teachers and students are discovering the very powerful effect that computer graphics can have on their art programs. Students may work in multiple types of media and, with a scanner or digital camera, import these materials into the computer for further manipulation and modification. Students who love photography may find new and exciting ways to express themselves by scanning their pictures and processing them with software.

Students can also output their work in a variety of ways: color or black-and-white, inkjet or laser, scaled down or blown up into poster-tiles. Each of these allows students and teachers to experiment and find the best way to present a given image.

The images that you see here represent a variety of projects and techniques. Some were done as backdrops for poems that the students composed. Others were the result of a high school 3D computer animation course that stresses creation of models and materials. Yet others were created by manipulating an image on the computer, printing it out for overhead projection, and then using this scene to create the final artwork.

Our jury included two high school teachers, an educational researcher who works with elementary and middle school students, and a professional computer animator. We hope you enjoy their selections.

Scott Lang
sigKIDS Art Guy

sigKIDS Jury

Heidi Dunphy
Independent Animator

Tim Comolli
South Burlington High School Imaging Lab

Maria Roussos
University of Illinois at Chicago

Chris Carey
Orange County Schools

Erin McCartney

Age 16

Digital Self-Portrait

Teacher
Kristy Higby
The Mercersburg Academy

Hardware
Macintosh computer
Digital camera
Overhead projector

Software
Adobe Photoshop

Neddly Maxime

Age 14

Carla

Teacher
Nancy Krebsbach
Maynard Evans High School

Hardware
Gateway 2000 computer

Software
Kinetix 3D Studio Max

A.J. Thieblot

Age 18

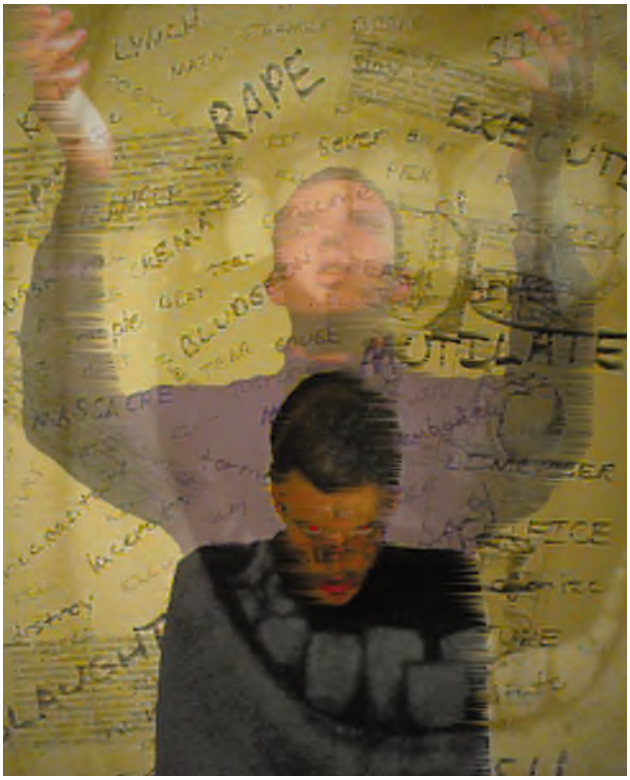

Digital Self-Portrait

Teacher
Kristy Higby
The Mercersburg Academy

Hardware
Macintosh computer
Digital camera
Overhead projector

Software
Adobe Photoshop

Matthew Teichman

Age 17

Bath Scene

Teacher
Kevin McTiernan
Academy for the Advancement of Science and Technology

Hardware
Hewlett-Packard Pentium Pro computer

Software
Kinetix 3D Studio Max
Adobe Photoshop

Jessica Yuan

Age 9

California Landscape

Teachers
Ramona Otto
Jane Shimotsu
The Mirman School

Hardware
Macintosh Power PC

Software
Fractal Design (Metacreations)
Dabbler
Hyperstudio

Daniel Sergile

Age 17

Split Personality

Teacher
Nancy Krebsbach
Maynard Evans High School

Hardware
Gateway 2000 computer

Software
Kinetix 3D Studio Max

Sarah Dodson

Age 9

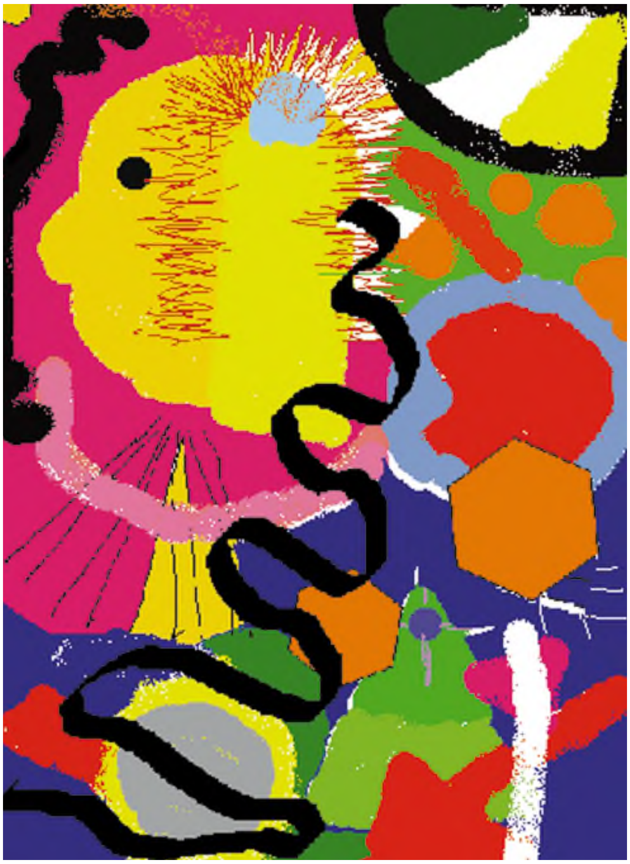

Untitled #87

Teacher
Jean L. Perry
Germantown Elementary School

Hardware
Macintosh computer

Software
ClarisWorks

Angel R. Espinoza

Age 15

My Apartment

Teacher
Nancy Krebsbach
Maynard Evans High School

Hardware
Gateway 2000 computer

Software
Corel Draw 7

Charlcie Legler

Age 18

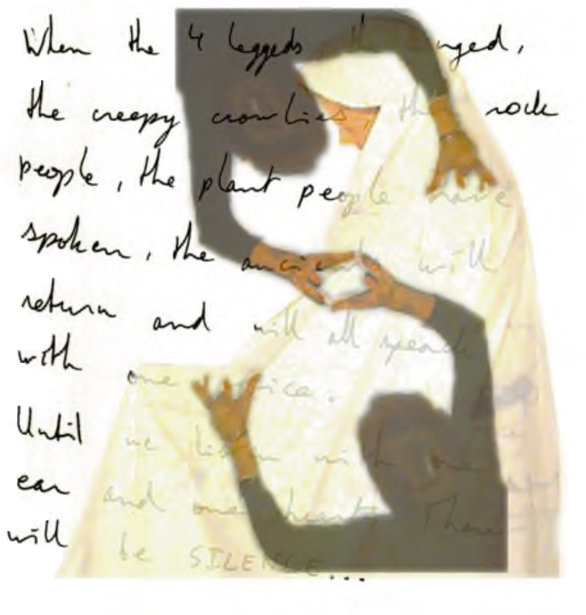

Digital Self-Portrait

Teacher
Kristy Higby
The Mercersburg Academy

Hardware
Macintosh computer
Digital camera
Overhead projector

Software
Adobe Photoshop

Matthew Teichman

Age 17

The Stregnar Civilization

Teacher
Kevin McTiernan
Academy for the Advancement of Science and Technology

Hardware
Hewlett-Packard Pentium Pro computer

Software
Kinetix 3D Studio Max
Adobe Photoshop

Laura Griffiths

Age 8

Trees

Teacher
Leanne Statland
The Mirman School

Hardware
Macintosh Power PC

Software
Fractal Design (Metacreations)
Dabbler
ClarisWorks

Joshua Hendle

Age 15

Life Dream

Teacher
Nancy Krebsbach
Maynard Evans High School

Hardware
Power Computing (Macintosh) computer

Software
Strata Studio Pro

Eddie Kang

Age 17

Digital Self-Portrait

Teacher
Kristy Higby
The Mercersburg Academy

Hardware
Macintosh computer
Digital camera
Overhead projector

Software
Adobe Photoshop

Eddie Kang

Age 17

Digital Self-Portrait

Teacher
Kristy Higby
The Mercersburg Academy

Hardware
Macintosh computer
Digital camera
Overhead projector

Software
Adobe Photoshop

Gizelle Mallillin Pera

Age 9

The Tree

Teacher
Leanne Statland
The Mirman School

Hardware
Macintosh Power PC

Software
Fractal Design (Metacreations)
Dabbler
ClarisWorks

Toujour Byrd

Age 15

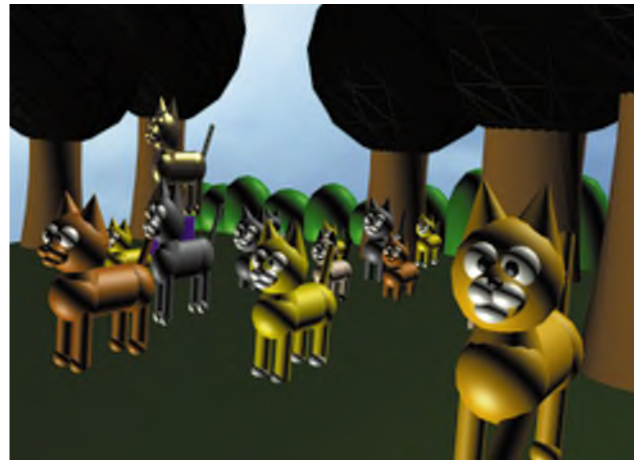

Alluromania

Teacher
Nancy Krebsbach
Maynard Evans High School

Hardware
Gateway 2000 computer

Software
Kinetix 3D Studio Max

Mackenzie Wahl

Age 8

I Saw a Tree

Teacher
Leanne Statland
The Mirman School

Hardware
Macintosh Power PC

Software
Fractal Design (Metacreations)
Dabbler
ClarisWorks

Dayana Ottenwalder, Eileen Lied, Juan Diaz, and Zulma Gomex

Age 17, Age 17, Age 17, and Age 18

Pyramid of Darkness

Teacher
Nancy Krebsbach
Maynard Evans High School

Hardware
Pentium 120 computer

Software
ClarisWorks

Chip Collier

Age 16

Untitled

Teacher
Nancy Krebsbach
Maynard Evans High School

Bettina Santo Domingo

Age 9

Moonlit Men

Teacher
Leanne Statland
The Mirman School

Hardware
Macintosh Power PC

Software
Hyperstudio

Alexandra Greene

Age 9

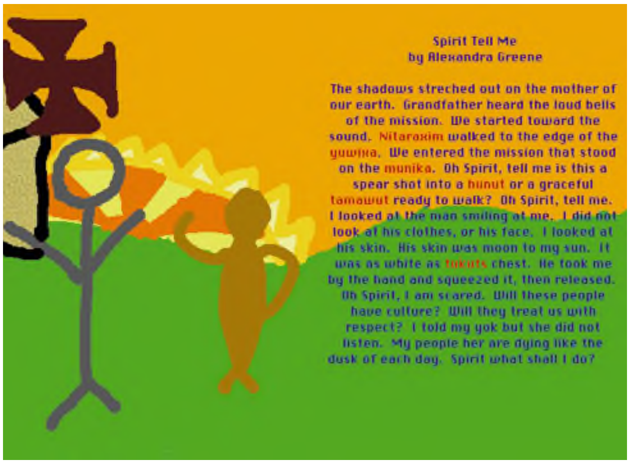

Spirit Tell Me

Teacher
Leanne Statland
The Mirman School

Hardware
Macintosh Power PC

Software
Hyperstudio

Jillian Banks

Age 9

What Does it Feel To Have Freedom?

Teacher
Leanne Statland
The Mirman School

Hardware
Macintosh Power PC

Software
Hyperstudio

Claire Abramowitz

Age 9

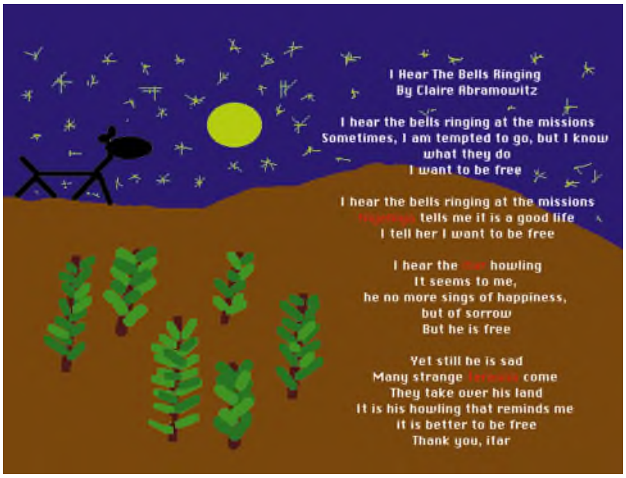

I Hear the Bells

Teacher
Leanne Statland
The Mirman School

Hardware
Macintosh Power PC

Software
Hyperstudio

David Halperin

Age 8

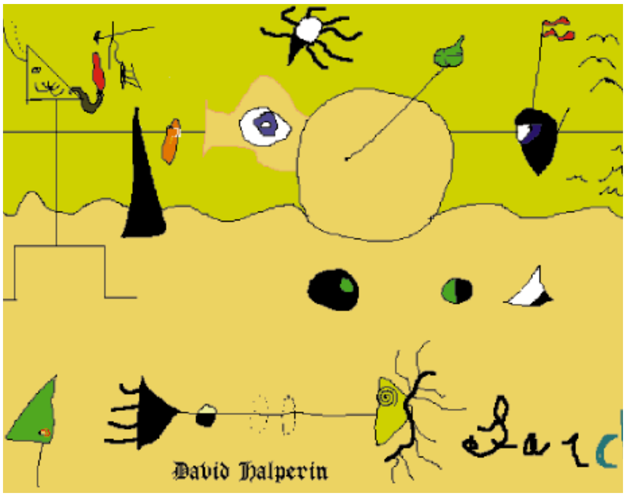

Spirits of Wonder

Teacher
Candace Corliss
The Mirman School

Hardware
Macintosh Power PC

Software
ClarisWorks and KidPix

Emily Rosenthal

Age 8

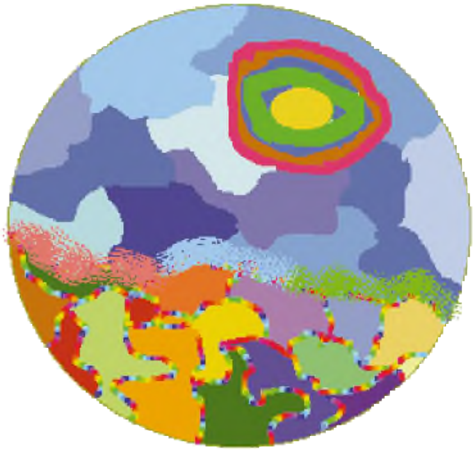

Circle of Life

Teacher
Candace Corliss
The Mirman School

Hardware
Macintosh Power PC

Software
ClarisWorks and KidPix

Elizabeth Crowson and Kinsey Harris

Age 9

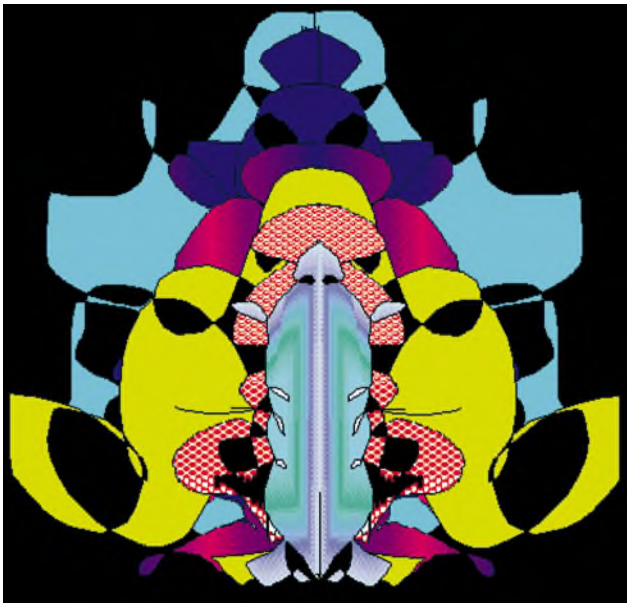

Untitled #11

Teacher
Jean L. Perry
Germantown Elementary School

Hardware
Macintosh computer

Software
ClarisWorks

Lam Nguyen

Age 18

Isolation 3.11

Teacher
Nancy Krebsbach
Maynard Evans High School

Hardware
Power Computing (Macintosh) computer

Software
KPT (Metacreations) Bryce 2

Vinh Nguyen

Age 18

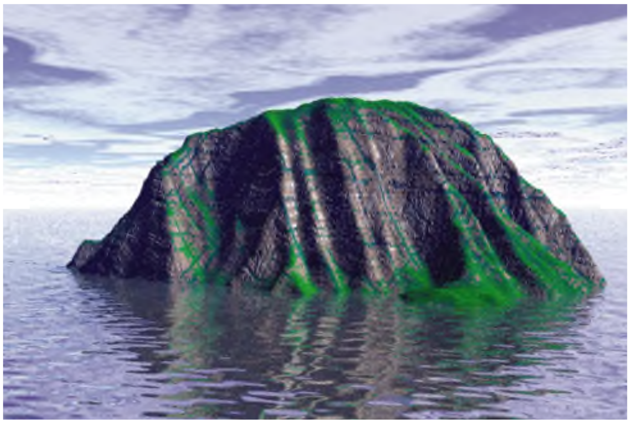

Pacific Island

Teacher
Nancy Krebsbach
Maynard Evans High School

Hardware
Pentium 90 computer

Software
KPT (Metacreations) Bryce 2

Natasha Spottiswoode

Age 6

The Mayflower Voyage

Teachers
Ellen Brown and Anita Tilley
The Mirman School

Hardware
Macintosh Power PC

Software
ClarisWorks and KidPix

Eunji Mah

Age 18

Digital Self-Portrait

Teacher
Kristy Higby
The Mercersburg Academy

Hardware
Macintosh computer
Digital camera
Overhead projector

Software
Adobe Photoshop

Conference Abstracts and Applications **sigKIDS / Community Outreach** 173

Panels

The Year of the Human

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

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

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

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

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

Celia Pearce
SIGGRAPH 98 Panels Chair

Visualization: The Hard Problems

Organizer
David Zeltzer
Sarnoff Corporation

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

Jock D. Mackinlay
Xerox PARC

Krzysztof Lenk
Dynamic Diagrams

Randall W. Simons
Sandia National Laboratories

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

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

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

So the hard problems remain:

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

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

David Zeltzer
Sarnoff Corporation

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

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

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

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

Ann M. Bisantz
State University of New York at Buffalo

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

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

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

Krzysztof Lenk
Dynamic Diagrams

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

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

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

Information Visualization: Using Vision to Think

Jock D. Mackinlay
Xerox PARC

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

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

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

Randall W. Simons
Sandia National Laboratories

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

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

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

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

Panelists
Dan Mapes
LEGO A/S, SPU-Darwin

Henry Sowizral
Sun Microsystems

Andries van Dam
Brown University

Dan Venolia
Cosmo Software

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Listen Up! Realtime Auditory Interfaces for the Real World

Moderators
Maribeth Back
Xerox PARC

Elizabeth Mynatt
Xerox PARC

Panelists
Perry R. Cook
Princeton University

Robin Bargar
Virtual Environments Group, NCSA

Peter B.L. Meijer
Philips Research Laboratories

Designing Interactive Auditory Displays

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

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

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

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

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

Audio Augmented Reality

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

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

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

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

Auditory Display Using Real-Time Sound Synthesis and Processing

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

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

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

Sound Authoring for Real-Time Synchronous Display in Immersive Environments

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

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

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

Cross-Modal Sensory Streams

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

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

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

Out of the Box: Toys Break The Screen Barrier

Moderator
Steve Schklair
Quantum Arts

Panelists
Christian Greuel
LEGO A/S, SPU-Darwin

Andy Rifkin
Mattel Media

Erick Strommen
Microsoft Corporation

Michael Patrick Johnson
Massachusetts Institute of Technology

Steve Sutyak
Hasbro Interactive

Steve Schklair
Quantum Arts
steve@quantumarts.com

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

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

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

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

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

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

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

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

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

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

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

Andy Rifkin
Mattel Media
rifkinan@mattel.com

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

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

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

Erick Strommen
Microsoft Corporation
erikstr@microsoft.com

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

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

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

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

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

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

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

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

Feature FX: Money Pit or Gold Mine?

Moderator
Patricia Rose Duignan
VFX Marketing
Management Consultant

Panelists
Scott Ross
Digital Domain

Jim Morris
Industrial Light & Magic

Carl Rosendahl
Pacific Data Images

Richard Hollander
Blue Sky | VIFX

Phil Tippett
Tippett Studio

Ray Feeney
RFX, INC.

Scott Ross
Digital Domain

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

Carl Rosendahl
Pacific Data Images

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

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

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

Ray Feeney
RFX, INC.

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

Phil Tippett
Tippett Studio

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

Jim Morris
Industrial Light & Magic

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

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

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

Richard Hollander
Blue Sky | VIFX

The bad part of the visual effects industry:

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

The good part of the visual effects industry:

It is fun.

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

Dis-Illusion of Life: Becoming a Digital Character Animator

Organizer
Barbara Mones-Hattal
Industrial Light & Magic

Moderator
Jacquelyn Ford Morie
Blue Sky | VIFX

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

Ken Perlin
New York University

James Sayers
Sheridan College

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

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

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

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

Endla Burrows
Industrial Light & Magic
endla@lucasdigital.com

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

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

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

Daniel Jeannette
Industrial Light & Magic
dj@lucasdigital.com

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

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

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

Peter Docter
Pixar Animation Studios
doc@pixar.com

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

Ken Perlin
New York University
perlin@nyu.edu

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

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

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

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

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

Human Factors in Virtual World Design: Psychological and Sociological Considerations

Moderator
Elizabeth Reid Steere
Royal Melbourne Institute of Technology

Panelists
Lynn Cherny
AT&T Research Labs

Mary Czerwinski
Microsoft Research

Tammy Knipp
Florida Atlantic University

Beth Kolko
University of Texas at Arlington

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

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

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

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

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

Representations and Agency in World Design

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

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

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

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

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

Beyond Designing For Usability

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

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

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

Peak Experiences In Virtual Environments: A Sudden Surge Of Meaning

Tammy Knipp
Florida Atlantic University
tknipp@fau.edu

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

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

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

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

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

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

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

Virtual Reality as Healing Art

Moderator
Galen Brandt
Virtual Healing

Panelists
Dorothy Strickland
Stetson University

Richard Satava
Yale University

Hugh Lusted
BioControl Systems, Inc.

Tom Riess
HMD Therapeutics

Rita Addison
Virtual Reality Artist/Consultant

Myron Krueger
Artificial Reality Corporation

Hunter Hoffman
Human Interface
Technology Lab

Galen Brandt
Virtual Healing
GalenB42@aol.com

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

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

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

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

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

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

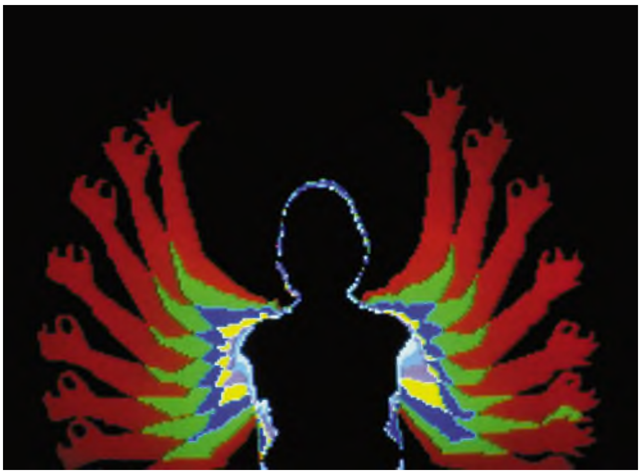

The healing art of Myron Krueger's VIDEOPLACE.

Virtual Reality as Healing Art

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

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

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

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

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

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

An autistic child practices street crossing on a virtual street.

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

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

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

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

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

Tom Riess
HMD Therapeutics
74244.1521@compuserve.com

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

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

"Virtual Freedom." Painting by Tom Riess.

Virtual Reality as Healing Art

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

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

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

Myron Krueger
Artificial Reality Corporation
myronkru@aol.com

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

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

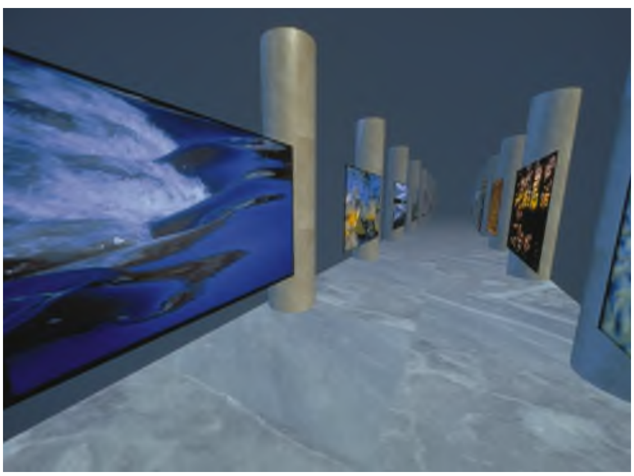

The virtual gallery in "DETOUR" showing Addison's photographs before her brain injury.

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

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

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

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

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

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

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

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

Moderator
Michael Harris
NCR Human Interface Technology Center

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

William T. Freeman
Mitsubishi Electric Research Laboratory

Hiroshi Ishii
Massachusetts Institute of Technology

Mark Lucente
IBM Research

Michael J. Sinclair
Microsoft Research

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

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

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

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

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

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

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

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

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

User Domains and Appropriate Design

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

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

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

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

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

Computers Looking at People

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

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

Tangible Media

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

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

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

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

Natural Interaction

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

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

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

HCI Through Creative Plagiarization

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

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

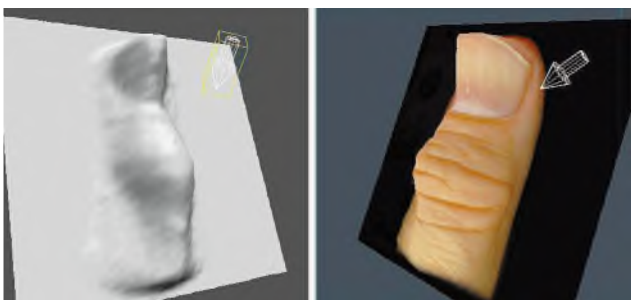

The Haptic Lens acquires 3D surfaces in real time, while applying a known force on the object or body part. (Michael J. Sinclair)

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

Moderator
Mel Slater
University College London

Panelists
Nat Durlach
Massachusetts Institute of Technology

Lawrence J. Hettinger
Logicon Technical Services, Inc.

Randy Pausch
Carnegie Mellon University

Dennis R. Proffitt
University of Virginia

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

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

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

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

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

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

Presence and Task Performance

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

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

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

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

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

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

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

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

Designing Successful Virtual Environment Applications

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

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

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

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

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

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

Getting People to Behave as if they are Present

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

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

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

Evoking Natural Perceptions of Simulated Scenes

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

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

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

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

The Sorcerer's Apprentice: Ubiquitous Computing and Graphics

Organizer
W. Bradford Paley
Digital Image Design Incorporated

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

Steve Feiner
Columbia University

Hiroshi Ishii
Massachusetts Institute of Technology

S. Joy Mountford
Interval Research Corporation

Steve Shafer
Microsoft Research

Mark Weiser
Xerox PARC

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The Sorcerer's Apprentice: Ubiquitous Computing and Graphics

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

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

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

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

Behavioral Modeling and Animation: Past, Present, and Future

Organizers
Demetri Terzopoulos
University of Toronto/Intel Corporation

Xiaoyuan Tu
Intel Corporation

Panelists
Kiran Joshi
Walt Disney Feature Animation

Ken Perlin
New York University

Craig Reynolds
DreamWorks Animation

Toby Simpson
Cyberlife Technologies Ltd.

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

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

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

The Design of Characters with Complex Behavior

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

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

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

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

Behavioral Animation in Disney Feature Films

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

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

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

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

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

Wildebeest stempede in "The Lion King."

AI Modeling for Behavioral Animation

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

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

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

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

Crowd in "The Hunchback of Notre Dame."

Artificial Fishes in a Digital Sea

Artificial Life in Home Entertainment

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

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

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

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

What are the Limits of Behavioral Modeling and Animation?

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

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

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

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

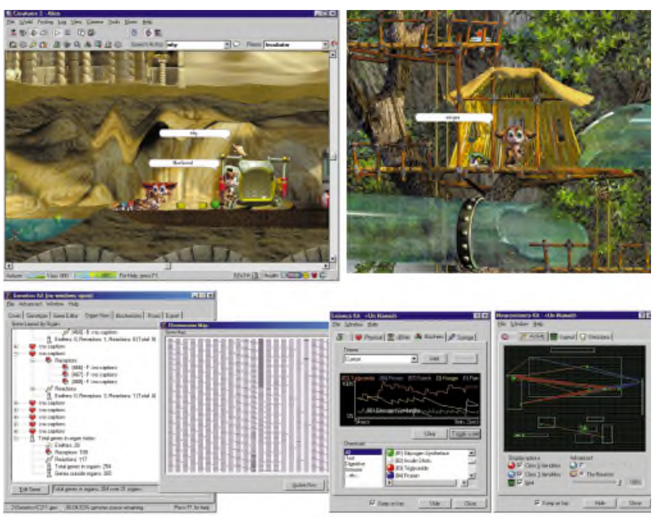

Creatures 2

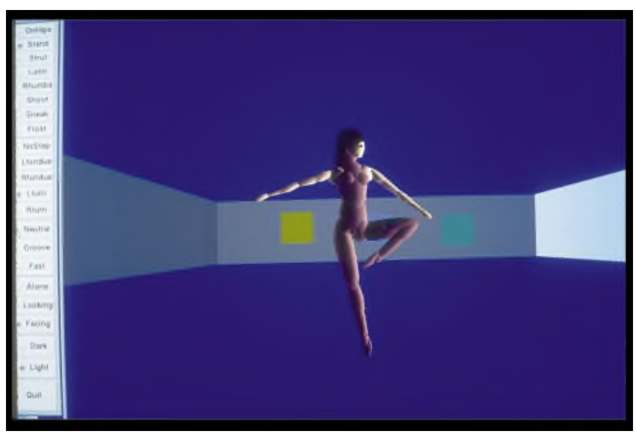

"Danse Interactif" by Ken Perlin, SIGGRAPH 94 Electronic Theater

Location-Based Entertainment: The Next Generation

Moderator

Randy Pausch
Carnegie Mellon University

Panelists

Trevor Bryant
Sony Development

Joe Garlington
Walt Disney Imagineering

Jon Snoddy
GameWorks

Jordan Weisman
FASA Interactive

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

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

The Meaning and Context of "Location"

Trevor Bryant
Sony Development
trevor@sre.sony.com

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

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

Shared Activity, Storytelling, and Creative Process

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

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

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

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

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

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

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

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

Jon Snoddy
GameWorks
jon@gameworks.com

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

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

Lessons Learned from Three Generations of "BattleTech"

Jordan Weisman
FASA Interactive
jweisman@fasainteractive.com

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

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

Sublime and Impossible Bodies

Moderator
Sara Diamond
Banff Centre for the Arts

Panelists
Jane Prophet
Slade School of Art

Joshua Portway
Realworld Studios

Catherine Richards
The University of Ottawa

Douglas MacLeod
WurcNet Inc.

Arlindo Ribeiro Machado Neto
Catholic University of São Paulo

Critical Art Ensemble (Steve Kurtz)

Ahasiw Maskegon-Iskwew
SOIL Digital Media Production Suite

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

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

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

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

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

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

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

Joshua Portway
Realworld Studios
josh@realworld.on.net

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Is Robust Geometry Possible?

Organizer
Kevin Weiler
Autodesk

Panelists
Tom Duff
Pixar Animation Studios

Steve Fortune
Bell Laboratories

Chris Hoffman
Cardinal Technology, Purdue University

Tom Peters
University of Connecticut

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

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

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

Topics:

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

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

Tom Duff
Pixar Animation Studios
td@pixar.com

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

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

Is Robust Geometry Possible?

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Kevin Weiler
Autodesk, Inc.
kjw@autodesk.com

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

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

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

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

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

Computer Vision in 3D Interactivity

Organizer
Mark Holler
Intel Corporation

Panelists
Ingrid Carlbom
Bell Laboratories, Lucent Technologies

Steven Feiner
Columbia University

George Robertson
Microsoft Research

Demetri Terzopoulos
University of Toronto / Intel Corporation

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

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

Audio-Visual Tracking for 3D User Interfaces

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

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

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

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

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

Steven Feiner
Columbia University
Feiner@cs.columbia.edu

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Ray Tracing and Radiosity: Ready for Production?

Organizer
Jacquelyn Ford Morie
Blue Sky | VIFX

Moderator
Richard Hollander
Blue Sky | VIFX

Panelists
Chris Wedge
Blue Sky | VIFX

Gonzalo Garramuno
Digital Domain

Grant Boucher
Station X Studios, LLC

Bob Powell
Rhythm & Hues Studios

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

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

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

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

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

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

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

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

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

Ray Tracing and Radiosity: Ready for Production?

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

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

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

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

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

Characters on the Internet: The Next Generation

Moderator
Steve DiPaola
Darwin Digital - San Francisco

Panelists
Barrett Fox
infoplasm

Athomas Goldberg
New York University

Mark S. Meadows
Construct

Celia Pearce
Celia Pearce & Friends

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Removing the Fourth Wall: Creating Interactive 3D Worlds

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

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

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

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

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

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

Interactive Dance Club

Introduction

The SIGGRAPH 98 Interactive Dance Club is a multi-participant interactive environment with real-time computer-generated imagery, lighting, and video, synchronized to dance club music (i.e. acid jazz, tribal, ambient, drum & bass).

Instead of dancing to prerecorded music and images, or passively watching a performance, members of the audience become participants.

Within interactive zones located throughout the club, participants influence music, lighting, projected computer graphics images and video. There are zones for single participants, dual participants, and groups of participants. Like sections in an orchestra, output from the interactive zones combines with a pre-defined basic rhythm to form the overall performance. Moving from zone to zone, participants experience different blends of musical and visual elements.

A sophisticated system of hardware and software keeps all the zones in sync while analyzing and filtering participant input, in order to deliver a musically coherent and visually satisfying experience. Feedback to the participants is designed to be immediate and responsive.

The Interactive Dance Club is designed to:

- Encourage and reward physical movement.
- Allow unencumbered interaction. No wires, no gloves, no goggles.
- Make learning to “work” the interactive zones intuitive and simple.
- Ensure that actions trigger meaningful responses in the context of the overall performance.
- Be simple, immediate, and fun!

The producers have been working to create an Interactive Dance Club that:

- Explores group and individual participation in the modulation of computer graphics, musical, and lighting elements while maintaining a musically satisfying and visually coherent whole.
- Demonstrates how fun it is for participants when they get their hands on the “steering wheel.”
- Creates an experience that amplifies the uniqueness of the individual and reveals the synergy of the group.

Ryan Ulyate
SIGGRAPH 98 Interactive Dance Club Chair
Synesthesia, LLC
+1.310.455.0585
+1.310.455.0595 fax
415 South Topanga Canyon Boulevard #210
Topanga, California 90290 USA
ryan@synesthesia-LLC.com

SIGGRAPH 98 Interactive Dance Club Committee

Interactive Dance Club Chair
Interactive Music Producer
Ryan Ulyate
Synesthesia, LLC

System Designer
David Bianciardi
Synesthesia, LLC

Hardware Integration
Dan Comins
Interactive Technologies

Computer Graphics Supervisor
Judith Crow
Digital Domain

Video System Designer
Tim Z. Falconer
Digital Domain

Project Manager
C.J. Flynn
OpsCenter Technologies, Inc.

Computer Graphics Software
Greg Hermanovic
Side Effects Software

Associate Lighting Designer
Andris Kasparovics
design one corporation

Video Display System Designer
Mike McLean
Hoffman Video Systems

Technical Director/Lighting Designer
Christien Methot
design one corporation

Interactive Technologies
Michael Moorhead
Hardware Integration

Computer Graphics Supervisor
Paul Simpson
Realise Studios

Coordinator
Kristen Stratton
Warner Bros. International Television Production

Design Consultant
Philip Vaughan
Vertex Productions

Interface Coordinator/Integrator
Peter Wyngaard
InterConnect of Ann Arbor, Inc.
Contributors

Computer Graphics Artists
Shannan Burkley
Matthew Butler
John Courte
Stewart Gordon
Craig Halperin
Caleb Howard
Carol Hunt
Alan Kapler
Franklin Londin
Kevin Mack
Martha Mack
Morris May
Rodney McFall
Scott Petill
Jarrett Smith
Hannah Walker
Petria Whelan

Audio Loops and Samples
Clilly Castiglia
Jerome Groove
Jonatas Manzolli
Mark Thorley

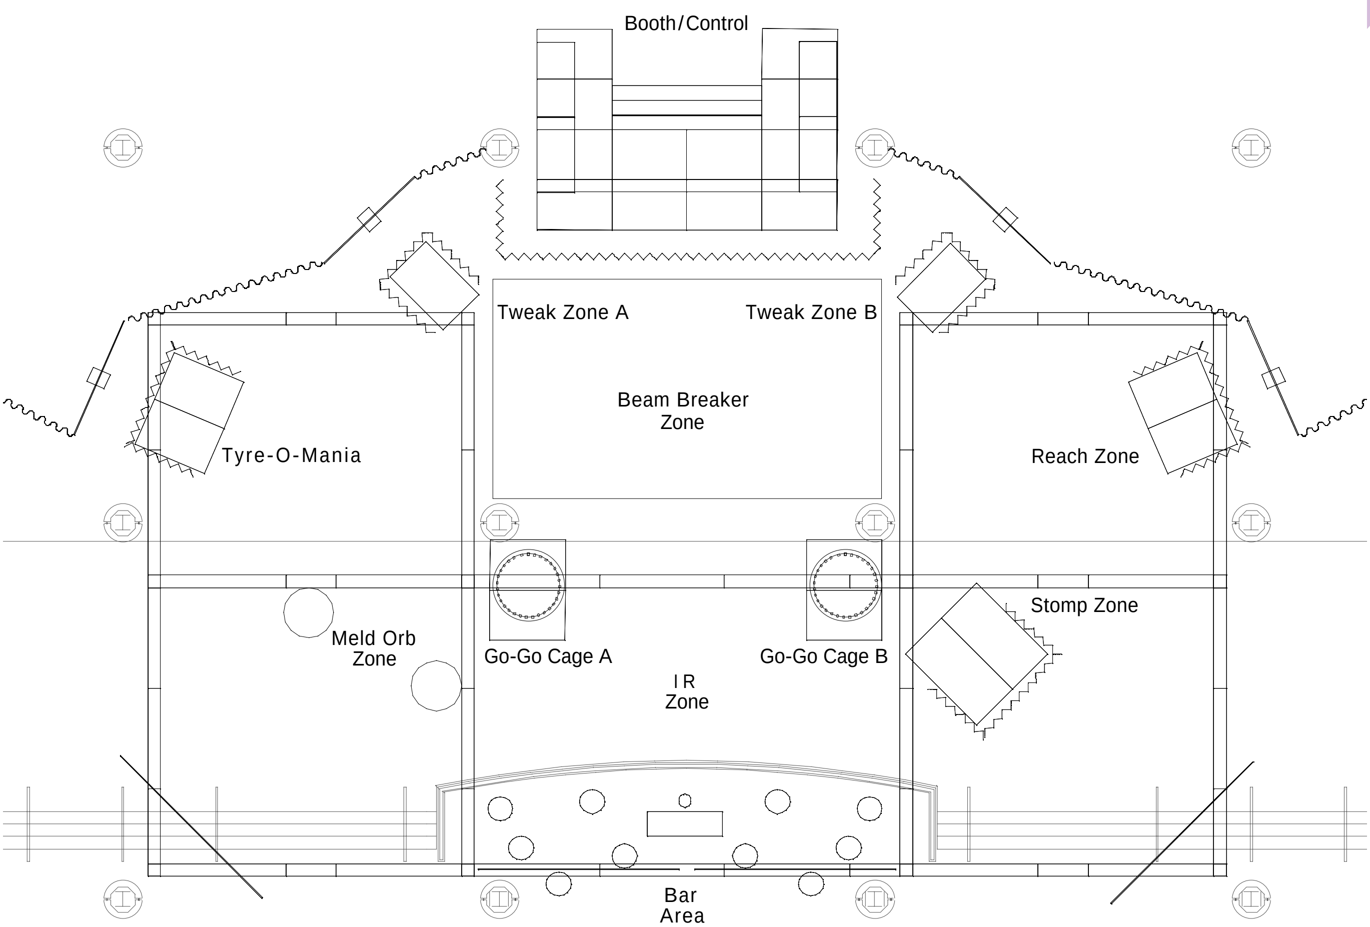

Themes

Several musical and visual elements can be modified by the participants in the various interactive zones. The content in all the zones is unified by four themes that define visual style, audio tempo, and tonality, and provide the context for four different overall experiences.

Are the Tiki Gods Angry?

Music style
Tribal

Example
Demon Priests

Image
Stuart Gordon / Paul Simpson

Awake In The Dream

Music style
Ambient

Example
Painting and Exploring Noise in Space

Image
Kevin Mack

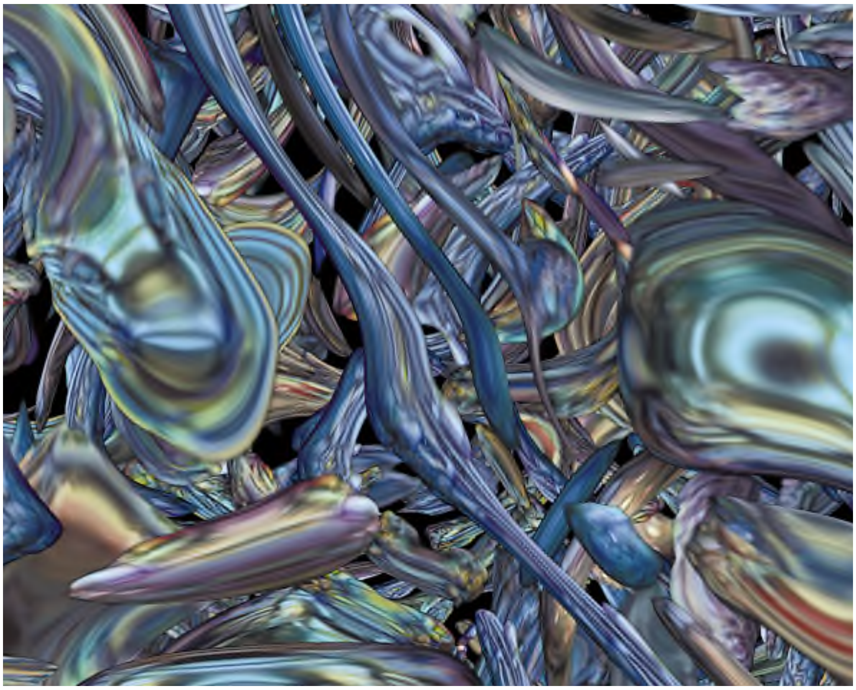

Millennium Fever

Music style
Drum & Bass

Example
Gearflow

Image
Alan Kapler

Yesterday's World of Tomorrow

Music style
Acid Jazz

Example
Moire Dancers

Image
Paul Simpson

Zones

Beam Breaker

The Beam Breaker zone consists of parallel light beams above people's heads on a dance floor with sensors that detect when a beam has been broken. Participants control aspects of the lighting and trigger musical phrases.

Infrared

Look for the hottest people on the dance floor. Literally! A state-of-the-art infrared video camera is focused on an area of the dance floor. The location of participants within this zone controls the surrounding lighting. Visual feedback is provided by projecting the infrared image on a large screen.

Stomp

Stomp is a two-person zone characterized by lots of physical motion. Participants interact by dancing and stomping on floor-mounted pads. Discrete pad hits and the level of activity in different areas within the zone control music and projected computer graphics.

Go-Go

Participants step up onto a platform two feet above the dance floor and enter a cone of light. Once inside the cone, they interact by extending their arms and breaking the surrounding light beam, casting a shadow on a circular array of sensors embedded in the platform below. This controls the music in the immediate vicinity.

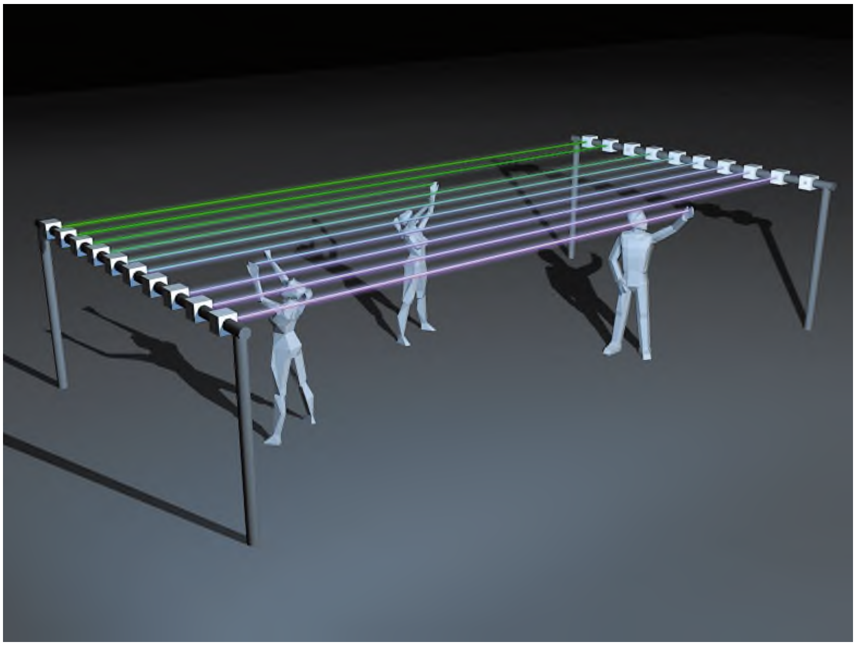

Beam Breaker   Rendering by Scott Petill

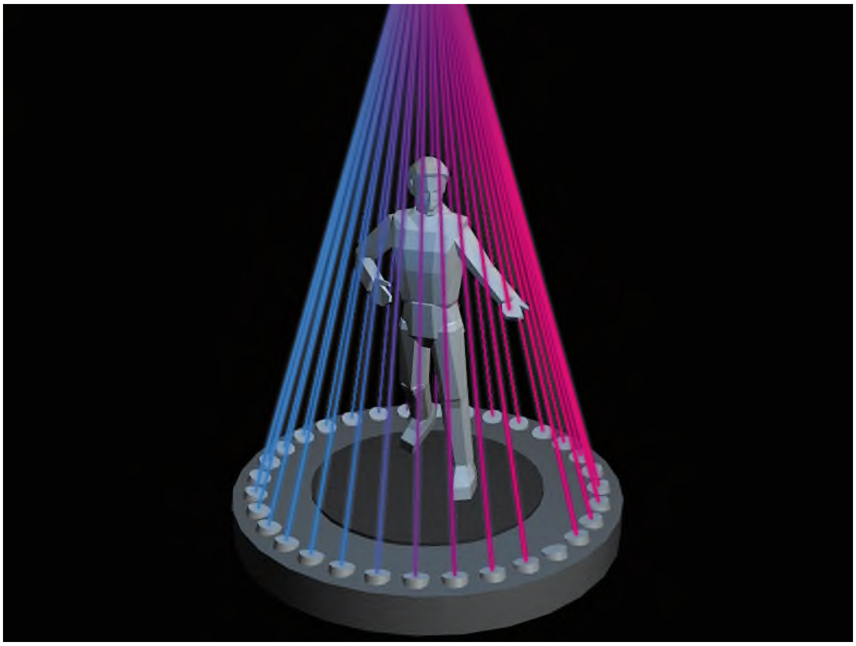

Go-Go   Rendering by Scott Petill

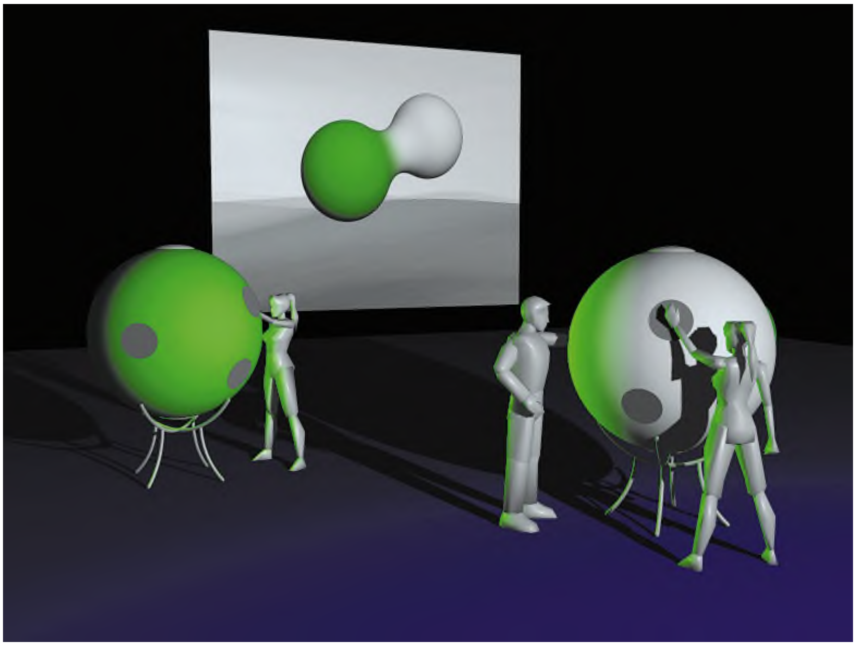

Meld Orbs   Rendering by Scott Petill

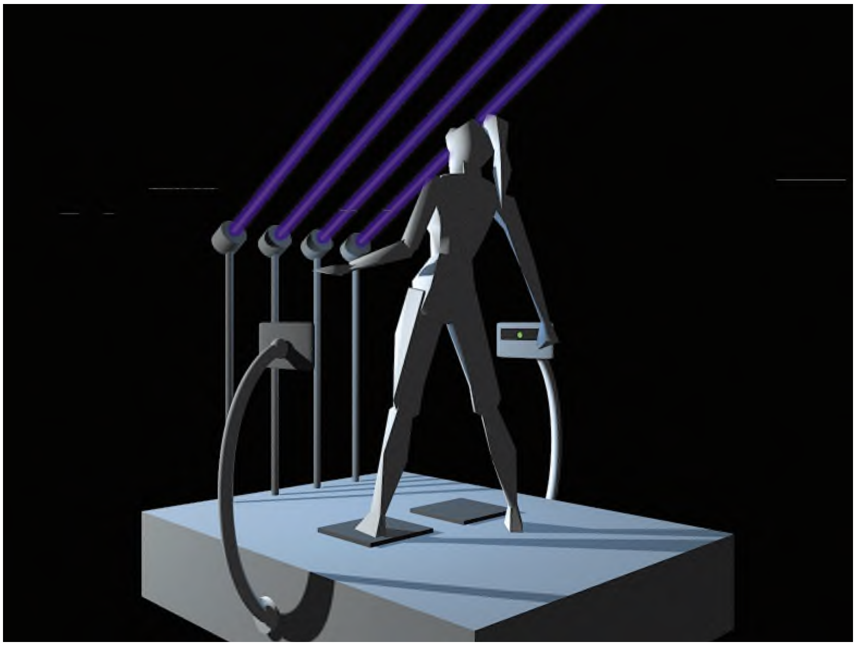

Tweak   Rendering by Scott Petill

Meld Orbs

The Meld Orbs are two four-foot-diameter spheres placed approximately 12 feet apart. Mounted on the surface of each Orb are nine near-field proximity sensors. Up to three participants interact with each Orb to affect sound and projected computer graphics imagery.

Reach

Reach is a one-person zone positioned on a raised platform. A video camera captures the participant's image, which is processed in real time, generating a computer graphics kaleidoscope. The participant must reach in order to touch four trigger pads positioned in a half-circle. Each pad controls an aspect of the kaleidoscope and sound in the immediate vicinity.

Tyre-O-Mania

Participants rotate two wheels mounted on a horizontal bar to affect projected computer graphics. A pair of foot pedals allows the participant to change which aspects of the computer graphics the wheels are controlling.

Tweak

Tweak is a one-person zone in which the participant interacts with two infrared proximity sensors and four overhead beam breakers. These interactions affect music and projected computer graphics.

Sketches

Technical

Sketches

Technical

Applications

Art, Design, Multimedia

Animation

Introduction

SIGGRAPH is a collage of animators, artists, designers, software developers, technicians, and researchers. In the midst of this diversity, the Sketches program is a forum in which all conference participants can become involved no matter what their background, training, or interests. The Sketches program is designed to promote discussion and free exchange of ideas by providing a smaller, more informal setting for presentations within the SIGGRAPH conference.

Sketches is a singular opportunity for animators, artists, and designers to present the technical aspects of their creations to the SIGGRAPH community. For end-users, Sketches is the place for presenting application of computer graphic and interactive techniques to solve real-world problems. For researchers, Sketches fulfills two functions: a conference-level platform for technical presentations of completed work and a forum for new, late-breaking, partially complete ("sketchy" work.

Sketches is five years old, still in its infancy, and it continues to evolve as it refines its role within the SIGGRAPH Conference. From Andrew Glassner initiated it at SIGGRAPH 94, Sketches has grown to become an active and thriving program. This year, the changes made to Sketches are subtle and have more to do with process than substance. Sketches contributions continue to be published as one-page extended abstracts, usually including one image. This format provides documentation of the ideas presented but does not preclude subsequent publication of the ideas in a more detailed format.

There are four categories of Sketches: Technical; Applications; Art, Design, and Multimedia; and Animation. The Art, Design, and Multimedia title has been expanded from SIGGRAPH 97 to include Multimedia in order to further encourage diversity of participation. Sketches received a total of 210 submissions this year; approximately half were in the Technical category. Each category had its own jury of experienced professionals, who evaluated submissions based solely on the quality of the submission. Of the 210 submissions, 86 were accepted. These numbers are similar to last year's.

Technical Sketches continue to be a source of interesting ideas, techniques, and innovations, including those concerned with: processing polygonal meshes, techniques in support of man-machine interaction, motion control, volume graphics, illumination, and rendering.

Applications Sketches cover various topics, including virtual reality, interactivity, and visualization.. One emphasis in Applications Sketches is on applying computer graphics to solve problems, from viewing a new hairstyle to inspects Michelangelo's Pieta. Three Applications Sketches present work that is on display in Emerging Technologies.

Art, Design, and Multimedia Sketches are dominated by artists using computer technology as an integral part of their creations. In many of these pieces, the interplay between human and machine creates a symbiotic artwork. In addition to these contributions, there are several sketches that present techniques for using digital processes to aid artistic design, bringing new and innovative ways of using the computer into the creative process.

Animation is once again an exciting Sketches category. It is an opportunity to hear practitioners discuss the details of techniques they have used to produce a particular piece of animation and to explain generic techniques used in support of animation production.

We have a wonderful Sketches program this year due to the efforts of many dedicated people. On behalf of the SIGGRAPH community, I'd like to express deep appreciation and thanks to the volunteers who served as category chairs and jury members. It is the dedication and efforts of these and other SIGGRAPH volunteers who make the SIGGRAPH experience possible. I would especially like to thank the Sketches' Administrative Assistant, Viki Dennis, for her outstanding work, and without whose help this job never would have been completed.

Enjoy!

Rick Parent

Sketches Chair
Rick Parent
The Ohio State University

Sketches Committee

Sketches Chair
Rick Parent
The Ohio State University

Administrative Assistant
Viki Dennis
The Ohio State Univeristy

Animation Subcommittee

Chair
Jonathon Luskin
Independent

Greg Griffith
Walt Disney Feature Animation

Ellen Poon
Industrial Light & Magic

Andrew Schmidt
Pixar

Applications Subcommittee

Chair
Marc Kessler
University of Michigan

Leo Catallo
Advanced Technology Laboratories

Gabriel Taubin
IBM

Jeff Vroom
Art Technology Group

Art, Design & Multimedia Subcommittee

Chair
Thecla Schiphorst
Emily Carr Institute of Art & Design

Sara Diamond
Banff Centre for the Arts

Diana Gromala
University of Washington

Patricia Levin
University of California at Irvine

Rebecca Allen
University of California at Los Angeles

Technical Subcommittee

Chair
David S. Ebert
University of Maryland Baltimore County

Andrew Glassner
Microsoft Research

John C. Hart
Washington State University

Jessica K. Hodgins
Georgia Institute of Technology

Christopher D. Shaw
University of Regina

Too Many Triangles

Wolfgang Seibold and Geoff Wyvill
Department of Computer Science
University of Otago
Box 56
Dunedin, New Zealand
geoff@rabbit.otago.ac.nz

You can halve the number of triangles used in many computer models with no loss of accuracy!

When we use triangles to approximate a curved surface, everyone seems to put the vertices on the surface. If they are close to the surface rather than on the surface, we can achieve better accuracy or use fewer triangles.

Suppose a regular polygon approximates a circle (Fig. 1). The error between the circle and the polygon is: $e = 1 - \cos(a)$ where $a = p/n$ and n is the number of sides. By moving the vertices, we can get an error: $d = (1 - \cos(a)) / (1 + \cos(a))$, reduced by a factor, $C = 1 / (1 - \cos(a))$. For large n, $C \approx 0.5$.

In 3D, the circle becomes a sphere and the polygon becomes a triangle mesh. We cannot build a regular mesh except in special cases, so we use Delaunay triangles1,2 to approximate the sphere. If the center of the circumcircle of a triangle lies inside the triangle, then the maximum error for that triangle occurs at this center (Case 1). Otherwise, the maximum error occurs at the mid-point of the longest side (Case 2). When we move the vertices, we can reduce the error by $C = 1 / (1 - \cos(a))$ where a is defined to be half the angle subtended at the center of the sphere by the diameter of the circumcircle (Case 1) or half the angle subtended by the longest side (Case 2). For large n, we halve the error as in 2D.

For small angles, $1 - \cos(a) \approx a2 / 2$, so the error is halved when the edge lengths are divided by 2. When we double the number of triangles, we halve the area of the average triangle, dividing its edge lengths by 2. So, when the vertices are constrained to lie on the surface, you need twice as many triangles for the same accuracy.

We cannot prove this for general surfaces, but it seems likely. If the triangles are small, then locally convex or locally concave areas will be covered by many triangles, and in those areas our theory applies.

But how do we decide where the vertices should be? Our algorithm finds the error of each edge as a vector. At each vertex, the total error, v, is the vector sum of the errors of edges meeting at that vertex. We shift each vertex by -kv, and find k by an iterative process that attempts to minimise v. We move the vertex in the direction of v because that works on the sphere.

Our test surface, Fig. 3, exhibits parabolic, hyperbolic, and elliptic points and therefore exercises the algorithm over a range of local shape variations. Our adaptive polygonization3 produces some poorly chosen vertices, so that the error sometimes increases even when a vertex is added. Fig. 4 shows the proportional decrease in maximum error when the vertices are moved off the surface as a function of the number of vertices. This proportion varies from 1.1 (an increase) down to 0.2. The average is about 0.7. This is larger than the 0.5 predicted, but we do not yet have the best off-surface placement of the vertices.

We cannot improve hand-built polygon models because there is no well-defined underlying surface that has been approximated. But parametric surfaces, where there is such an ideal, are better represented by triangles whose vertices do not lie on the surface.

References
1 Musin, O.R. Properties of the Delaunay triangulation. Proceedings of the Thirteenth Annual Symposium on Computational Geometry, Nice, France, June 1997, 424-426.
2 Wyvill, G. and McRobie, D. Adaptive polygonization of cao en surfaces. Insight through Computer Graphics Proceedings of CG International 194 World Scientific, 1996, 250-256.
3 Seibold, W. and Wyvill, G. Towards an understanding of surfaces through polygonization. Proceedings of CGI 198, in press.

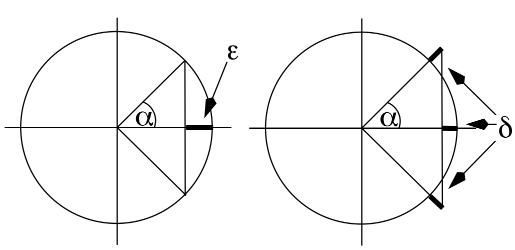

Figure 1 Polygon approximation to circle

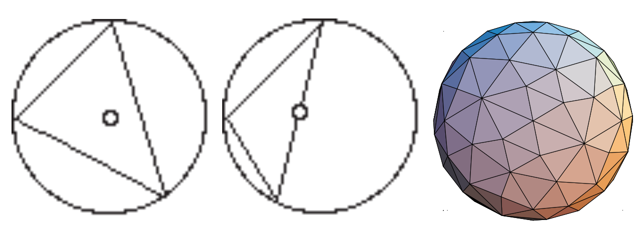

Figure 2. a) Max error Case 1, b) Max error Case 2, c) Delaunay sphere

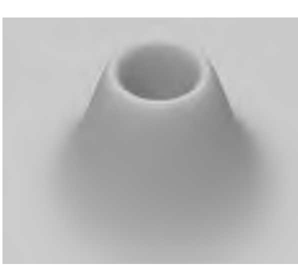

Figure 3 The test surface

Figure 4 Proportional reduction in maximum error for 2 to 1,000 polygons

Multiresolution of Arbitrary Triangular Meshes

Wei Xu and Don Fussell
Texas Institute of Computational and Applied Mathematics and Department of Computer Sciences
The University of Texas at Austin
Austin, Texas 78712 USA

+1.512.475.8632
+1.512.471.8694 fax
wei@ticam.utexas.edu

1 Motivation

Current multiresolution and topology-preserving representations of the topology of triangular meshes can be classified into two major categories: lossless and lossy. Lossless methods usually refine meshes progressively with vertex-adding techniques.[2,3] Such approaches can reconstruct the original mesh perfectly from a simple initial mesh, but they do not provide good parametrizations relating surface levels. Lossy methods apply regular (quaternary) surface subdivision schemes to base meshes to approximate irregular meshes.[1,4] Although they can not recover the original topology (connectivity) exactly, the regular subdivision schemes enable us to construct 2D wavelets on triangular meshes so that geometry and color information of meshes can be expressed hierarchically and efficiently.

To combine the advantages of these two different approaches, we are developing a new topology representation for arbitrary triangular meshes. Our method expresses the topology of meshes hierarchically and losslessly, and enables us to construct 2D wavelets on the arbitrary triangular meshes.

2 The Approach

The key theory (general subdivision theorem) supporting our approach is follows (see 6 for detail). A very general class of triangular meshes, called normal triangular meshes, can be represented as subdivision trees. Each interior node of the tree, representing a triangle, is subdivided using one of four elementary subdivision operations (shown in Fig.1). According to the theorem, a normal triangular mesh can be reconstructed by applying a sequence of elementary subdivisions to an initial single triangle.

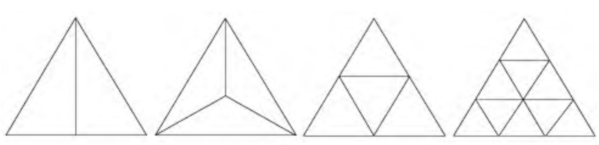

Figure 1 Four elementary subdivision schemes

The necessary and sufficient conditions for a mesh to be normal are simple and easy to verify.[6] We also present some methods to convert abnormal meshes into normal meshes at a small additional cost. These methods still allow us to reconstruct the initial input mesh losslessly.

A subdivision tree representation creates a sequence of nested spaces for the topology. These spaces give us a natural way to parameterize surfaces between different levels, thus making it possible to construct surface wavelets on the given meshes. Using the lifting scheme of 5, we construct wavelet-based representations of arbitrary triangular meshes with piecewise triangular Bézier patches as the basis functions. For piecewise linear functions, our wavelets are a natural extension of wavelets generated by hat functions. Unlike existing methods, our wavelets have closed forms for high-order cases, although this restricts us to using Bézier rather than B-spline surfaces for our wavelets, and thus orders of continuity higher than C^0 between patches can only be achieved with inter-patch geometric constraints as is usual with Bézier patches. Of course, C^0 continuity is all that is required for many graphics applications employing polygonal models.

Our surfaces allow a representation of approximations to the input surface at various levels of detail, adaptively selectable in different regions, as well as complete lossless recovery of the input mesh. Such properties are important for mesh compression as well as interactive surface interpolating and editing.

References

1 Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M., and Stuetzle, W. Multiresolution analysis of arbitrary meshes. In Proc. of SIGGRAPH 95, 173-182, 1995.
2 Hoppe, H. Progressive meshes. Proc. of SIGGRAPH 96, 99-108, 1996.
3 Li, J.K., Li, L., and Kuo, C.C. Progressive compression of 3D graphics models. Proc. of IEEE International Conf. on Multimedia Computing and System, Ottawa, 1997.
4 Schröder, P. and Sweldens, W. Spherical wavelets: Efficiently representing functions on the sphere. In Proc. of SIGGRAPH 95}, 161-172, 1995.
5 Sweldens, W. The lifting scheme: A custom-design construction of biorthogonal wavelets. Tech. Report, 1994:7, Industrial Mathematics Institute, Dept. of Math., Univ. of S. Carolina, 1994.
6 Xu, W., Fussell, D. Subdivision tree representation of arbitrary triangular meshes. Technical Report No. 9803, CS Dept, Univ. of Texas at Austin, Jan., 1998.

Geometric Reconstruction with Anisotropic Alpha-Shapes

Michael Capps and Marek Teichmann
MIT Lab for Computer Science
545 Technology Square, NE43-242
Cambridge, Massachusetts 02139 USA
+1.617.253.2598
+1.617.253.6652
marekt@graphics.lcs.mit.edu

Generation of a three-dimensional model from an unorganized set of points is an active area of research in computer graphics. Such point sets come from a number of common sources, such as range data from three-dimensional scanning hardware, implicit surface, and medical imaging.

The notion of α-shapes provides an elegant mathematical framework for extracting the geometric structure of a set of points in three dimensions.[2] Briefly, the α-shape is a set of triangles and tetrahedra that is a subset of the Delaunay triangulation of the input point set. The theory of α-shapes provides a method for obtaining a surface by selecting a subset of the triangles in the triangulation.

A triangle is in the α-shape if one of the radii of the following spheres is at most α: the spheres circumscribing the triangle's vertices and the two nearest neighboring points, and the sphere circumscribing the triangle. While this definition gives good results for point sets of roughly uniform density with large separation between surfaces, this definition is clearly not optimal for non-uniform point sets, or for surfaces that are separated by a distance less than their sampling density. For these, there exists no value of α that includes all desired triangles and deletes all undesired triangles. See the first image in the figure for an example of the best possible surface obtained using α-shapes.

We propose two extensions to alleviate these problems in the case where normal information is available (or estimated as in Hoppe et al.[3]) at each point. These extensions allow reconstruction from a larger class of point sets:

1 Anisotropic scaling: we allow the spheres to vary in shape, and change the triangulation accordingly.

A fundamental contribution of this work is incremental retriangulation based on a user-specified factor τ, for the influence of the anisotropy. The spherical forbidden region, which depends on three or four points, is locally deformed along the local average normal direction d. It is compressed along d if the point normal and the normals of the triangle align well. It is stretched otherwise, decreasing the likelihood that this triangle will be selected for the α-shape. This in effect varies the local metric tensor.[1]

The amount of deformation depends on the normal correlation and is multiplied by an interactively adjustable parameter τ. The local normal direction d is multiplied by τ, so the user has direct control of the anisotropy; $\tau = 1$ generates the standard sphere. The user modulates τ and α to better find the desired triangulation of the input point set.

This deformation also affects the Delaunay triangulation of the points. A tetrahedron is part of the Delaunay triangulation of a set of points if no points lie in its circumsphere. This sphere is also altered by the above deformations. After a change of τ, the triangulation is incrementally recomputed to satisfy the new local criteria using Delaunay flipping,[2] a standard method for computing Delaunay triangulations. The deformations are actually implemented by scaling the points along the direction d by the appropriate amount.

2 Density scaling: we vary the value of α depending on the local point density.

We estimate the local density at each point by

$$\delta(p) = \sum_{q \in P} 1 - \frac{d(q,p)}{\lambda} \quad \forall q \text{ such that } d(q,p) < \lambda,$$

where λ is the constant radius of the local neighborhood and $d(x, y)$ is the Euclidean distance function.

Given the three vertices of a potential triangle t, we let

$$\delta(t) = \frac{\delta(a)+\delta(b)+\delta(c)}{3\mu}$$

(μ is the average density) and set the local value of α to

$$\frac{\alpha}{\delta(t)^{\sigma}}$$ where σ is another user-adjustable parameter.

We see in the second image of the figure that density-scaling α deletes all triangles that connect the higher-density ball to the lower-density cylinders nearby. In the third image, we add normal-induced anisotropy, which eliminates the unwanted triangles across the interstice.

To implement our system, we used the Delaunay triangulation package written by Ernst Mucke.[4] The implementation is sufficiently fast for interactively varying the τ and σ parameters (and retriangulating) for point sets of several hundred thousand points, after an initial precomputation of the Delaunay triangulation.

References
1 Frank J. Bossen and Paul S. Heckbert. A pliant method for anisotropic mesh generation. In Proc. 5th International Meshing Roundtable, 63-74. Also Sand. Report 96-2301.
2 H. Edelsbrunner and E. P. Mucke. Three-dimensional alpha shapes. ACM Trans. Graph., 13(1):43-72, January 1994.
3 Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. In Edwin E. Catmull, editor, Computer Graphics (SIGGRAPH 92 Proceedings), volume 26, 71-78, July 1992.
4 Ernst Mucke. Shapes and implementations in three-dimensional geometry. PhD thesis, University of Illinois at Urbana-Champaign, 1993.

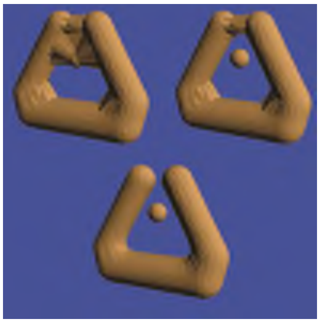

Converting Sets of Polygons to Manifold Surfaces by Cutting and Stitching

Andre Gueziec, Gabriel Taubin, Francis Lazarus, William Horn
IBM T.J. Watson Research Center
P.O. Box 704
Yorktown Heights, New York 10598 USA
{gueziec,taubin,francis,hornwp}@watson.ibm.com

A manifold polygonal surface is such that the neighborhood of every vertex can be continuously deformed to a disk (to a half disk at the boundary). This corresponds to an intuitive "surface" definition, as opposed to an arbitrary collection of polygons. Topological degeneracies can occur by design choice (e.g., vertex merging to avoid duplicating coordinates, or polygon reduction tools), or they can be produced by incorrect algorithms for building surfaces (e.g., iso-surfaces, triangulation of scattered points) or by correct algorithms containing software bugs, etc.

Concrete examples of algorithms that can fail on input containing topological singularities are: surface subdivision, surface simplification, surface compression, progressive transmission of geometry, etc. Possible approaches include: modifying algorithms to handle non-manifold input; locating the source of errors in modeling or CAD packages, and lobbying (and hoping) for such errors to be corrected; and developing methods to correct the input. The first approach is application-dependent and probably requires re-defining objectives (beyond just accepting non-manifold input: e.g., a number of surface-simplification methods accept non-manifolds, but often they introduce many more degeneracies than those originally present.) The second approach has little short-term impact and may not be a complete solution in the long term as well (there will always be software bugs).

We have chosen the third approach and developed an automated method for correcting topological singularities. The algorithm manipulates the polygon vertex indices (surface topology) and essentially ignores vertex coordinates (surface geometry). This differentiates our method from most of the previous work. Also, we do not assume that the method works with solids. We assume that the topology of the surface is already built for the most part. However, we provide a stitching method ("edge snapping") to help build the topology from disconnected polygons. Except for the optional stitching, the algorithm has a linear complexity in the number of vertices, edges, and faces, and requires no floating point operation.

Our algorithm comprises two steps: cutting and stitching. Cutting consists of disconnecting the surface along a collection of marked edges or vertices. Multiple copies of vertices are created and assigned new indices. Vertex indices in faces are modified to refer to the proper copy. We mark singular edges and vertices (an edge is singular if more than two faces are incident to it; a singular vertex is such that several edge-connected fans of triangles are incident to it. We propose two different methods for cutting: a local method and a global method. The local method operates only on marked vertices and endpoints of marked edges, by counting the number of (unmarked) edge-connected sets of faces incident to a vertex and by duplicating the vertex and assigning a different copy of that vertex to each connected set. The global method operates on all the faces and vertices of the surface at once, by first breaking all connections between faces and later joining adjacent faces that share an unmarked edge. The global method is more appropriate when there is a large number of topological singularities to correct. The local method is more efficient when there are only few singular elements in a generally correct topology.

Several manifolds can be mapped to the original non-manifold by identifying vertices. To reduce the number of vertices, holes, or components, the cutting operation is followed by stitching. Stitching consists of taking two boundary edges and identifying them. We propose two different stitching strategies called "edge pinching" and "edge snapping." Edge pinching consists of stitching adjacent edges on one boundary where a cut was previously made. Edge snapping consists of stitching edges whose endpoints are within an epsilon distance. Stitching can be a delicate operation: singular edges can be created when stitched multiple times. We avoid this by testing that after each proposed edge stitch, none of the affected vertices becomes singular.

To provide a real-world example, we consider a polygonal model of the space ship Enterprise with 12,539 triangles and 15,011 vertices. Figure1 shows a global view, where disconnected surface components are painted with different colors. We discovered 594 singular edges and 1,878 singular vertices. Figure 2 shows a detail with singular edges painted in red, boundary edges in green, and regular edges in yellow. After removing degenerate triangles (Figure 3), there are 435 singular edges and 1,689 singular vertices left. After cutting (Figure 4) and edge snapping (Figure 5), there are 12,552 triangles and 7,429 vertices. This example is a good advocate for automated correction methods: asking a user to decide on how to correct the surface locally 1,800 times seems impractical.

References
Murali, T.M. and Funkhouser, T.A. Consistent solid and boundary representations from arbitrary polygonal data. 1997 ACM Symposium on Interactive 3D Graphics, 155-161.

Barequet, G. and Kumar, S. Repairing CAD models. IEEE Visualization 1997.

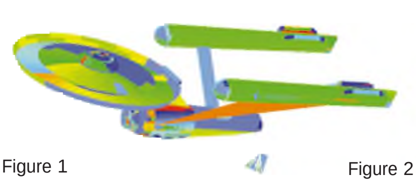

Figure 1

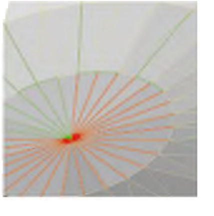

Figure 2

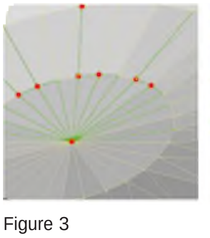

Figure 3

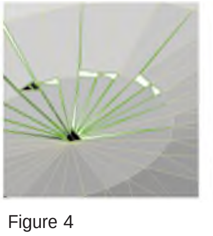

Figure 4

Figure 5

Collision Detection Framework Using Model Simplification

Tiow-Seng Tan, Ket-Fah Chong, and Kok-Lim Low
National University of Singapore
Department of Information Systems & Computer Science
Lower Kent Ridge Road
Singapore 119260
tants@iscs.nus.edu.sg

Though there have been significant advances in constructing good bounding volume hierarchies for collision detection (see, for example, Gottschalk, et al. OBBTree: A Hierarchical Structure for Rapid Interference Detection. SIGGRAPH 96, 171-179), there remain many controversial design issues, such as top-down versus bottom-up, and the choice of bounding volume (Klosowski et al. Efficient Collision Detection Using Bounding Volume Hierarchies of k-DOPs, State University of New York, Stony Brook, Dept. Applied Math. and Statistics, July 1997). This project proposes a new framework to resolve some of these design issues, using output from model simplification. It differs from the previous work of, for example, Klosowski et al. in looking beyond constructing different types of (possibly computationally expensive) bounding volumes to better understand the importance of shapes of objects in collision detection.

Our proposed framework uses output from simplification in the construction of bounding volume hierarchies. The basic idea is to use clues from simplified models to partition the given polygonal model into components with tight bounding volumes. A simplified model S contains elements such as points, edges, and triangles derived from the original model. We obtain a partition of these elements based on an equivalence relation, where element x and y are related if and only if x is connected to y by triangles in S. From parts of the partition, we then obtain components where each component contains those triangles of the original model that are simplified to elements in a same part. This process can be recursively applied, each time with a simplified model of higher level of detail, to obtain sub-components of components. All these components and sub-components can be arranged into a few topmost levels of the bounding volume hierarchy, and the remaining levels of the hierarchy can be computed using some other standard bounding volume hierarchy algorithms.

The above process exploits the fact that the set of triangles, edges, and points in a simplified model shows a sketch of the object, and its vertices are also more uniformly distributed compared to the original model. As such, computation to derive a bounding volume hierarchy with simplified model can be (1) effective due to its "shape" information captured by its essential elements, (2) unbiased for uniformly distributed vertices, and (3) faster due to the smaller number of vertices in the simplified model. The overall construction has the unique flavor of bottom-up in that triangles are distributed into various components, followed by top-down in that triangles are partitioned as in, for example, the RAPID [Gottschalk et al.]. On the other hand, the success of the framework also depends on the simplification that produces drastically simplified models that show good sketches of the given models. Equally important, the simplified models should be consistent across different levels of detail, since they place constraints on the recursive partitioning of components into sub-components. Our floating-cell simplification algorithm (Low & Tan. Model Simplification Using Vertex-Clustering. Proc. 1997 Symp. on Interactive 3D Graphics, 75-81) meets the above requirements and is adapted for our current implementation.

Preliminary experimental results show that the framework can improve existing collision detection systems in the construction of bounding volumes with little preprocessing overhead and without user intervention. For most 3D interactive applications, simplification cannot be considered as overhead, as both simplification and collision detection are necessary processes.

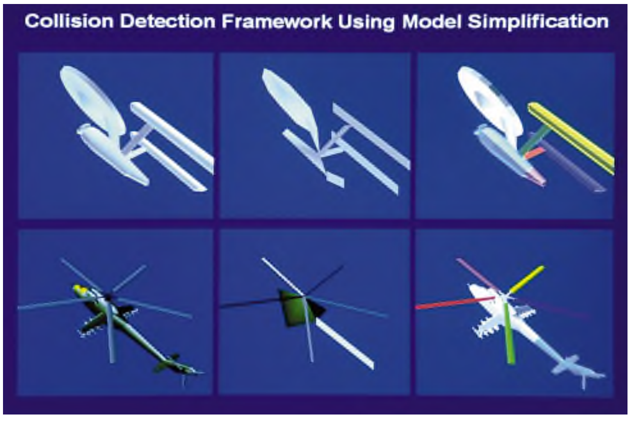

Optimizing Stereo Video Formats for Projection Based Virtual Reality

Gary Lindahl, Tom DeFanti, Dan Sandin, Greg Dawe, and Maxine Brown

Electronic Visualization Laboratory
University of Illinois at Chicago
851 South Morgan Street, Room 1120
Chicago, Illinois 60607-7053 USA

lindahl@evl.uic.edu
{tom|dan|dawe|maxine}@evl.uic.edu
www.evl.uic.edu

Many of today's high-end, projection-based virtual reality (VR) systems generate frame-interleaved video for stereoscopic equipment. Integrating LCD shutter glasses, high-end graphics computers, and large-scale projectors is difficult. This paper identifies problems discovered as part of our VR efforts at the Electronic Visualization Laboratory (EVL) and offers solutions.

Common Problems

There are four problems commonly encountered when integrating the SGI InfiniteReality (IR) graphics system,[2] Electrohome Marquee 8500 projector, and StereoGraphics CrystalEyewear. The problems are: phosphor decay lag or ghosting colors, stereo separation or left- and right-eye phase lock, video genlock or video vertical phase lock, and video interference and reflection.

Phosphor Decay Lag (Ghosting Colors)

A CRT projection system is made up of red, green, and blue phosphor tubes; the light from each tube combines to create full-color images. The common green phosphor tube has a slow decay rate that affects good stereo viewing separation (i.e., the green components of full-color images don't degrade as quickly as the red and blue components, so the user still sees the green component of previous images as images are updated; this is "ghosting"). The red tube also has a slow decay rate, but is acceptable for most people. To fix this lag problem, order a projector with a fast-decay P43 green phosphor tube.

Stereo Separation (Left- and Right-Eye Phase Lock)

Stereo separation, the ability to synchronize stereo images on a single screen, was a problem only recently corrected with the introduction of the IR. To see this problem, display the pattern, shown in Figure 1(A), in stereo. Wearing CrystalEyewear glasses, you will notice the defect, illustrated in Figure 1(B), at the top of the video frame: at the top edge of the video frame, the left and right-eye images are swapped. SGI corrected this in software by adding more pixel lines to the vertical back porch of the video frame, positioning the swapped left/right eye image above the frame. A disadvantage of this solution is that the signal bandwidth exceeds the projector display bandwidth and decreases the possible number of pixels that can be viewed on the screen; they now appear in the hidden portion of the video frame.

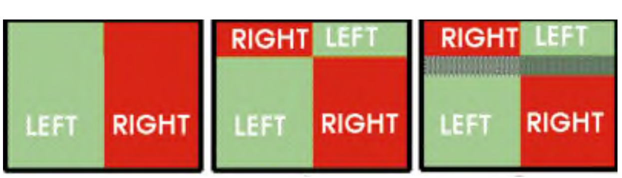

Figure 1 Video stereo display with (A) correct separation, (B) incorrect separation, and (C) stereo separation drifts. Image courtesy of Electronic Visualization Laboratory, University of Illinois at Chicago, and Pyramid Systems, Inc., Southfield, Michigan.

To correct for pixels off-screen, SGI adjusts the actual stereo sync pulse position relative to the video vertical-frame retrace by adding a function to the Video Format Compiler (VFC).[2] This code (Figure 2) makes all pixels visible and uses the VFC option to adjust stereo separation.

```
Line Based {
int linesHigh = [user defined];
int linesLow = [user defined];
signal "FIELD" initial state = low;
Transition Lines Range "FIELD" = linesHigh to linesHigh high at
BeginTime;
Transition Lines Range "FIELD" = linesLow to linesLow low at
BeginTime;
 }
```

Figure 2 The source code needed to change the stereo sync pulse location.

Video Genlock (Video Vertical Phase Lock)

Video genlock synchronizes the display of images from multiple video sources onto multiple screens. The CAVE uses multiple video sources, one for each wall of the CAVE, which require genlocking. Each video source "locks" to the master video source that supplies the stereo sync output. Without locking to the master source, a horizontal dark bar will appear and stereo separation will drift, as shown in Figure 1(C), when seen through the CrystalEyewear glasses. An SGI hardware solution is to connect the genlock sync input with the horizontal sync output of the master video source. A software solution is to configure the video output format of the slave source so it genlocks to the TTL or nominal video sync levels of the master video output. The software fix is usually the hardest since one has to develop a software-defined rule set that is compatible with the design of the video hardware.

Video Interference and Reflection

Ripple-like defects (noise "interference" in the projector) or smearing (an echo-like "reflection" in which images fade and overlap) may occur in the video image frame. The ripple-like defect is caused by the projector not having enough retrace time to stabilize in order to scan the video frame cleanly. The VFC option lets us add pixels to the video format definition. Adding more pixels to the horizontal front porch gives the projector more time to retrace and stabilize. Smearing is corrected by using high-bandwidth-grade video cable or a video extender over fiber (if longer cable length is required).

Conclusion

EVL works with equipment manufacturers to help identify and correct problems such as these as newer generations of computer and display technologies become available for use with projection-based VR systems like the CAVE.

References

1 C. Cruz-Neira, D.J. Sandin, T.A. DeFanti. Surround-screen projection based virtual reality: The design and implementation of the CAVE. Proceedings of ACM SIGGRAPH 93, 135-142.

2 J.S. Montrym, D.R. Baum, D.L. Dignam, and C.J. Migdal. InfiniteReality: A real-time graphics system. Proceedings of ACM SIGGRAPH 97, 293-302.

Acknowledgments

EVL receives major funding from the National Science Foundation (NSF), the Defense Advanced Research Projects Agency, and the US Department of Energy; specifically NSF awards CDA-9303433, CDA-9512272, NCR-9712283, CDA-9720351, and the NSF ASC PACI program. We also acknowledge Silicon Graphics, Inc. and Advanced Network and Services for their support.

3D Painting for Non-Photorealistic Rendering

Daniel Teece
Regent Court
211 Portobello Street
Sheffield S1 4DP United Kingdom.
dan@dcs.shef.ac.uk
www.dcs.shef.ac.uk/~dan/

Session: Colliding, Caving & Coloring

In recent years, the field of non-photorealistic rendering has seen a transition from a predominance of simple 2D post-processors to 3D renderers. Within this spectrum, a combination of interactive and automatic systems is evident, but to date there have been few fully interactive 3D systems.

The 3D Expressive Painter is an interactive, expressive 3D painting and rendering system intended for use in production of hand-drawn natural-media-style animations. By using some imported scene geometry, created in a general-purpose modeller, the system allows interactive placement of paint strokes in 3D. The user is able to rotate the scene, select a suitable viewing angle, and paint from the desired direction, with the depth values for the stroke positions calculated from the imported scene. High-level effects can be applied to objects within the scene (for example, tracing determined "key edges" with pencil strokes or fillig surfaces with a solid colour).

Though it inherits many of the features of a typical 2D painting package, such as multiple-level undo and palette-selection facilities, the 3D Expressive Painter produces "3D paintings," as opposed to 2D images, by applying the drawing operations to solid surfaces rather than a flat canvas. Furthermore, the 3D Expressive Painter differs from 3D-texture-painting systems in that it does not use reverse texture mapping to create the painted scene. The artistic effects interactively defined by the user are largely view-independent, based in three-dimensional space, and multi-resolution in nature.

Once completed, these 3D paintings can be used to produce static renderings from any viewpoint at any resolution, and also as the basis for generating animations. After providing a camera-motion path, each frame in the animation is rendered by applying perspective transformations to the paint strokes to obtain the 2D projection of the painted scene. Basing all of the functionality in 3D helps animated output considerably. Achievement of a hand-drawn look using the random variation of strokes is usually performed in screen space, which can cause "line wobble" and "boiling" artifacts in animations, where the displacement for a given line or point varies wildly from one frame to the next. By using a combination of screen-space and world-space perturbation, the 3D Expressive Painter gives animators more control over the level of frame-to-frame coherence.

In addition to many of the normal tools expected from a painting package, several new features are provided, some of which use the 3D information in the imported model to aid the painting process. Tone Sourcing uses light-shading calculations at the scene-painting position to give the mark placed at that point an approximate tonal value. Object Selection gives the user the option of concentrating painting effects on certain parts of the 3D scene. Colour Referencing allows the artist to use colours from an imported source image in the composition. Also, 3D paintings can be stored independently of scene geometry, meaning that the painting can be saved, retrieved, and viewed without the original model.

Visibility, often a problematic issue for non-photorealistic rendering, is implemented in a choice of different modes, the most useful being Depth Testing and Solid Occlusion. Depth Testing uses the z-order of the strokes to determine visibility at the pixel level, whereas Solid Occlusion uses ray-intersection testing to eliminate brush marks hidden from view. These schemes have comparative strengths and drawbacks, and, with the provision of several visibility determination modes, the user can select whichever works well with the scene in question.

In the 3D Expressive Painter, a non-photorealistic look is achieved by avoiding reliance on the standard rendering pipeline and using new expressive drawing primitives. In this, the ideas of visual noise, artistic approximation and economy of line are incorporated to give a hand-drawn appearance to the final imagery.

This research was sponsored by LightWork Design Ltd.

Visually Representing Multi-Valued Scientific Data Using Concepts from Painting

David H. Laidlaw, David Kremers, Eric T. Ahrens, Matthew J. Avalos
Caltech Computer Graphics Lab
Mail Stop 139-74
Pasadena, California 91125 USA
{dhl,davidkremers,eta,avalos}@gg.caltech.edu
www.gg.caltech.edu/~dhl

Vector-valued and tensor-valued images are rich sources of information about many physical phenomena. Visually representing these images so that they can be understood is a challenge, however, because they contain so many inter-related components,[2,5] all of which must be represented simultaneously and intuitively.

We have applied concepts from oil painting to the display and understanding of multi-valued scientific data. For example, painters map components of a scene onto visual characteristics of brush strokes by varying their size, shape, color, texture, opacity, direction, and placement. They also convey information by stroke locations relative to one another. Using these variations in appearance, we can simultaneously represent many components of multi-valued data and show relationships among them. Painters also build up an image with multiple layers of paint, where each layer represents and encapsulates some components of the data. The lowest layer, or underpainting, often roughs out the form of the painting. Layers can be semi-transparent or sparse, and thus be built up without obscuring one another. We combined these techniques to create an interactive computer graphics system[4,7] for experimenting with visual representations of 2D images of vector- and tensor-valued data.

Artistically motivated rules guided our choice of brush characteristics and layering to represent components of the data.[6,8] Our examples are of three different data types: diffusion-tensor data showing the pathology of a mouse disease in the spinal cord,[1] 2D vector and tensor measurements of flow over an airfoil,[3] and six-valued magnetic resonance imaging data showing the embryonic development of the mouse brain. The images are effective because they display many data values simultaneously, they qualitatively represent the underlying phenomena intuitively and geometrically, and they emphasize different data values to different degrees, leading a viewer through the temporal process of understanding the relationships among them.

Support
NSF (CCR-96-19649); NIDA, NIMH, NSF (Human Brain Project); NSF (ASC-89-20219); Beckman Institute; Apple, DEC, HP, IBM.

References

1 E. T. Ahrens, D. H. Laidlaw, C. Readhead, C. F. Brosnan, S. E. Fraser, and R. E. Jacobs. MR microscopy of transgenic mice that spontaneously acquire experimental allergic encephalomyelitis. Magnetic Resonance in Medicine, (1998). (to appear)
2 T. Delmarcelle and L. Hesselink. Visualizing second-order tensor fields with hyperstream lines. IEEE Computer Graphics and Applications, 13 (1993), 25--33.
3 G. G. Gornowicz, Continuous-field image-correlation velocimetry and its application to unsteady flow over an airfoil. Eng. thesis, California Institute of Technology, 1997.
4 P. E. Haeberli, Paint by numbers: Abstract image representations. In Computer Graphics (SIGGRAPH 90 Proceedings), F. Baskett, ed., vol. 24, Aug. 1990, 207--214.
5 L. Hesselink, F. H. Post, and J. J. van Wijk. Research issues in vector and tensor field visualization. IEEE Computer Graphics and Applications, 14 (1994), 76-79.
6 G. Kreutz, Problem solving for oil painters. Watson-Guptill Publications, 1986.
7 B. J. Meier, Painterly rendering for animation. In SIGGRAPH 96 Conference Proceedings, H. Rushmeier, ed., Annual Conference Series, ACM SIGGRAPH, Addison Wesley, Aug. 1996, 477-484.
8 E. R. Tufte, The visual display of quantitative information. Graphics Press, Cheshire, Connecticut, 1983.

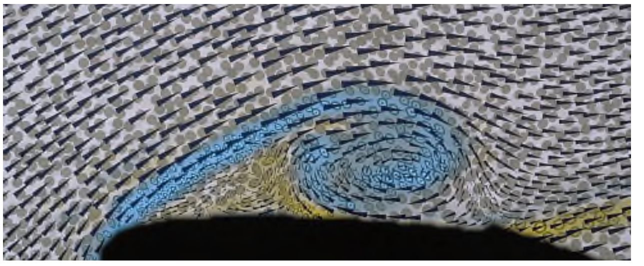

Figure 1 2-D flow field visualization. The image simultaneously displays the velocity (two values), the vorticity (one value), and the deformation-rate tensor (three values). The values are encoded, respectively, in the size and orientation of a layer of wedges, in a color base layer, and in the size, shape, orientation, opacity, and texture of a layer of ellipses.

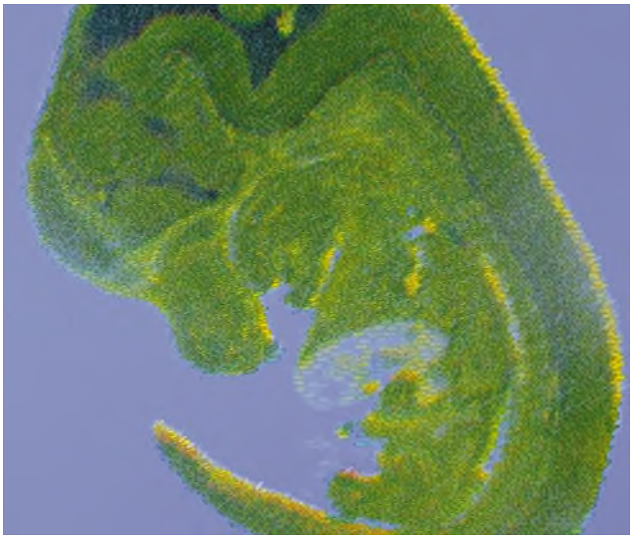

Figure 2 Simultaneous display of six-valued MR data. Four of the component data values are encoded in the color and transparency of an underpainting. The size and orientation of a layer of small elliptical outlines encode the other two values.

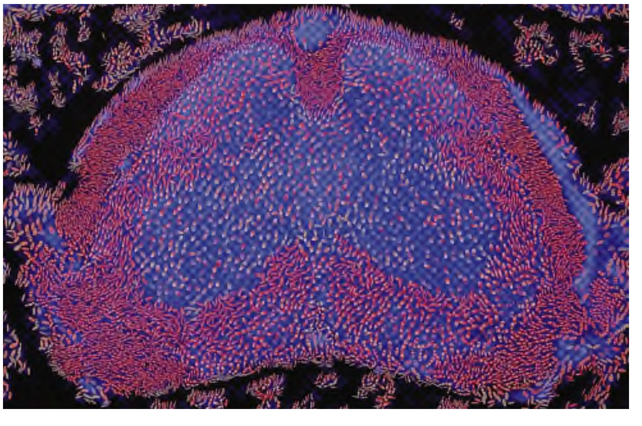

Figure 3 Layers of brush strokes displaying one slice of a second-order tensor field (six values) and two spatially correlated scalar fields (two values) over 3D. Eight interrelated variables are shown. This image is composited from four layers. Component data values are represented in the lightness of a purple underpainting; a transparent grid pattern composited over that; the direction, shape, color, and transparency of a layer of elliptical strokes over that; and the frequency of a texture applied to the elliptical strokes.

Star Cursors in Content Space: Abstractions of People and Places

P.J. Rankin
Philips Research Laboratories
Redhill, Surrey RH1 5HA United Kingdom
rankin; clive; jack@prl.research.philips.com

C. v Heerden, J. Mama, L. Nikolovska, R. den Otter, J. Rutgers
Philips Design
Building HWD-3, P.O. Box 218
5600 MD Eindhoven, The Netherlands
l.nikolovska; r.otter; j.rutgers@design.philips.com

Virtual multi-user worlds modeled on a literal representation of users and physical environments are currently the vogue in computer graphics.[1] Our assumption is that abstract representations are more expressive for visualising many domains,[2] offering new types of "super-natural" social navigation.[3] This sketch presents design concepts for an information environment, "ContentSpace," inhabited by abstract representations of users in the form of "StarCursors."

Exterior: The Outer Space and User Presence

ContentSpace is constructed from "content walls," which not only focus the topic domains of local zones, but also set their social properties such as seclusion, genuineness, or access for communication, shopping, work, or edutainment. Content walls respond to the gaze or proximity of users. The interactions between people and place are always two-way, zones mediating and contextualizing the interpersonal exchanges they enfold. Users are represented by a personalised avatar, a StarCursor. Its heart, body, eye, gaze beam, and aura act as multimedia channels, connected to the user profile for personal disclosures, communication or action (c.f.4).

The Inner Space

A novel graphical metaphor for their inner world enables users to verify their personal profile and control its context-dependent disclosure. StarCursors can be assigned classified interests, demographics, roles, tastes, etc. The interest profile filters the perception of other users and the environment. Default disclosures to others employ stereotype "masks" over facets of this data, according to pre-defined contexts such as job seeking, working, shopping, or romancing. In a specific zone, the user can therefore trust what will and will not be revealed. Body space between two cursors automatically moderates disclosures between two cursors and can be expanded or contracted to suit friends or predators. A body language is proposed, communicating aspects of identity via emission of multimedia objects, triggered by the gaze and user profile of the beholder. Data gathering on activities and locales further augment a user's profile, thus highlighting issues of trust.

Observations and Future Work

A first prototype of the interaction between two StarCursors and with a video wall was coded in Java and VRML, interfaced via the External Authoring Interface. VRML's Level of Detail construct proved convenient, although weaknesses in current browser support for subtle audio soundscapes, overlapping translucency, and light beams were exposed.

Unifying the interface designs to support both inner and outer worlds in one abstract metaphor proves richly productive and provocative. Connections among many research fields appear and a host of new interface design challenges emerge. Anticipated future work is in further development of secure and adequate representation of the inner world, identity disclosure processes, and human studies of aspects of trust and usability.

References

1 Damer, B. Avatars! Peachpit Press, Berkeley, 1998.

2 Donath, J. S. Inhabiting the virtual city. PhD thesis, Massachusetts Institute of Technology, February 1997.

3 Dieberger, A. Supporting social navigation on the World Wide Web. Int. J. Human-Computer Studies, 46, 805-825, 1997.

4 Greenhalgh, C. and Benford, S. MASSIVE: A collaborative virtual environment for tele-conferencing. Trans. ACM Computer-Human Interfaces, 2(3), 239-261, September 1995.

5 Platform for Privacy Preferences Project, WC3, see www.w3.org/Privacy/

Acknowledgements

This work builds on earlier concepts created with R, vd Haar and T.J. Everett. H.S. Hudson, Z. Jetha and A.T. Markettos provided VRML prototyping.

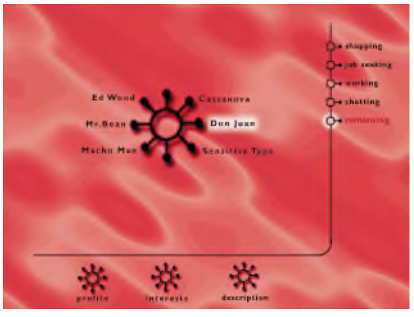

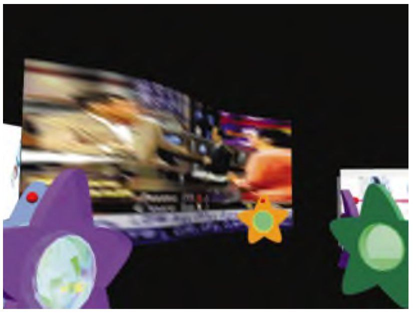

A 3D Stereo Window System for Virtual Environments

Michael A. Wingfield
The Mitre Corporation
202 Burlington Road
Bedford, Massachusetts 01730 USA
mwing@mitre.org

We are familiar with the 2D windows we use every day on our personal computers. These windows partition the desktop into overlapping areas that can be associated with different applications.[1] Similarly, the virtual environment consists of a workspace that extends the desktop metaphor into the 3D domain. By extension, we would like to provide a way of partitioning the 3D workspace (deskspace) into overlapping spaces that can be associated with independent scenes, objects, and applications. This sketch describes the implementation of a 3D stereo window system that enables a user to define mutually independent spaces within a virtual environment.

Design Goals

The following design goals were used to drive the specification of a 3D stereo window system. The same user interaction protocols that are used in the deskspace would be carried over to the interaction with objects in each window space. These interactions involve creating and direct manipulation of objects with datagloves in 3D space. Objects in one space would not have any relationship to objects in another space. One should be able to grab objects in one window space and move them seamlessly to another window space and have scale changes occur automatically. Trackballs should be associated with a window to permit the user to fly through the scene in a window independently of other windows or the workspace. Each window space should have its own independent and variable scale. Lights within a window should illuminate only objects in that window. Finally, windows should be viewed as independent objects in their own right in the workspace, so they can be grabbed, moved around, resized, and iconified.

Implementation

Because users view a virtual environment in stereo 3D, it follows that the view looking into each window must also be in stereo 3D. Since the window may be set to any orientation within the deskspace, it is important that careful attention be paid to the projection mathematics and location of the center of projection (COP).[2] The COP for each window must be the same as the deskspace and must be located at the user's viewpoint.

Our implementation uses a technique of stereo texture mapping to obtain the image in the window. This involves generation of a left-right image pair using two associated projection matrices and rendering into an off-screen buffer which is then translated into a texture using standard OpenGL calls. The window face polygon is drawn with the texture component corresponding to the left or right eye view. A texture size of 256 by 256 pixels for each of the left and right views has been found to be a reasonable size when trading off performance versus fidelity.

A scene consisting of a meaningful arrangement of graphical objects can displayed either in the deskspace or in any window. We used the concept of a domain as an attribute of each object that identifies the space within which it resides. Domains effectively partition the total number of objects into sets, within which only certain operations are permitted. For example, only objects in the same domain can be grouped together. Collision detection is among objects in the same domain. When an object is removed from one window, it is assigned to the deskspace domain, and when it is moved into another window, it is assigned to the domain associated with that window. The user interface for this operation involves reaching into one window, grabbing an object, moving it out through the window face, and moving it into a second window.

The concept of an active window on a workstation screen provides focus for both keyboard and mouse. In some implementations, the location of the cursor identifies the focus, while others use a window-highlighting feature. We have extended this concept to 3D windows by allowing the user to select the window that is the current focus. The active window then has all of the user-interface functions associated with it. For example, a joystick would control objects in the active window. Selection of another window would attach the joystick to that window. The active window concept is very useful in such operations as scene loading, scene deleting, grabbing and manipulating objects with datagloves, and scene scale changes. Currently, a window is made active by touching the window frame with a dataglove pointing gesture.

Summary

Extension of the workstation 2D window metaphor into 3D stereo windows in a virtual environment provides an additional tool for displaying and interacting with visual information. Multiple, independent scenes can be displayed in the virtual environment. and a consistent user interface capability permits manipulation of objects in any space. We are continuing to experiment with different user interface paradigms appropriate for this way of viewing virtual environments.

References

1 James D. Foley, A. van Dam, S. Feiner, and J. Hughes. Computer Graphics Principles and Practice. Addison-Wesley Publishing Company, Second Edition, 1990, 439-469.
2 Ibid., 229-283.

Visorama: A Complete Virtual Panorama System

Andre Matos, Luiz Velho, Jonas Gomes,
Andre Parente, Heloisa Siffert
Rua Duque Estrada 31 Apt. 606
Gavea
Rio de Janeiro, RJ CEP 22451-090 Brazil
amatos@visgraf.impa.br

Session: Reality ++

Introduction

Among the first image-based rendering systems were the virtual panorama systems. In a panorama, the user can look freely around a point in the virtual environment but cannot move continuously. Several such currently available systems differ in a number of ways, but they all have a few common limitations. Some do not provide a natural and immersive interaction with the panorama-based virtual environment. We present the Visorama System, which uses new software and hardware components to enable an immersive interaction with panoramas. In addition, it has a set of authoring tools that allow creation of panoramas and specification of the environment's multimedia structure. This is the first complete system that provides all these components.

Immersive Panoramas

Most existing systems fail to provide an immersive interaction with panoramas because the viewing direction is not correlated to the user's head motion, but is manipulated using the mouse. Several devices have been developed that could be used to solve these problems, such as head-mounted displays or the BOOM.1 These devices, however, are not appropriate for panorama-based environments because they provide more degrees of freedom than the panorama system. As a result, there might be a loss of synchronization between the user's visual and physical senses, which will lead to loss of immersibility and possibly to motion sickness. To avoid these problems, a visualization device to be used with panoramas must only have two degrees of freedom, for changing the viewing direction, which match those in the panorama system. This type of limited interaction has several advantages: the navigation paths can be easily specified during authoring, 2D multimedia data can be realistically integrated into the environment, and no 3D information is needed for collision detection, which is important in image-based environments. This is a typical case where "less is more." Limiting the user's navigation freedom makes it easier to create environments with more complex behaviors.

The Visorama Hardware Device

As part of the Visorama System, we developed a hardware device that solves the problems mentioned above. Figure 1a illustrates an artistic rendering of the Visorama observation device, and Figure 1b shows the first prototype. The device uses a binocular display by N-Vision Inc. to show the image generated by the panorama system. This display is attached to a support base that can rotate around vertical and horizontal axes, which have high-resolution sensors (5,000 positions) that together capture the current viewing orientation. In addition, three buttons allow the control of zoom angle and generation of discrete events. The sensors and buttons are sampled at 60 HZ, and their values are sent to the multimedia platform, where the output image is generated by the system's software. This form of direct manipulation of the viewing parameters provides a natural interface for virtual panoramas, as can be seen in Figure 2.

The Visorama Software

Image-generation software for immersive devices has requirements that are not satisfied by the existing virtual panorama systems. In particular, it must guarantee that there will never be any flickering or any latency between user actions and their visual feedback, which would result in a loss of immersibility. To satisfy this requirement, we developed a virtual panorama viewing engine that is compatible with OpenGL and uses hardware-implemented texture mapping to warp the panoramic image for visualization. As the amount of texture memory is usually not sufficient for a scene represented by panoramas with multiple levels of detail (we use a FireGL 4000 card with 16MB of texture memory), we designed an algorithm to limit the amount of data loaded into this memory per frame, avoiding latency. The image is broken into small tiles (e.g., 64x64), and the algorithm maintains an ordering of the estimated probability for each tile's visualization. The ones with higher probability are loaded into texture memory, one per frame, so that when tiles are needed for visualization they are already loaded.

Authoring and Applications

In addition to the visualization components, we are also developing a new set of authoring tools specifically for panorama-based virtual environments that allow the multimedia structure of the virtual environment to be specified using an event-based authoring language. Events are generated as the user navigates through the environment, rotating and zooming with the virtual camera. These operations determine which areas of the panorama are visualized and at which level of detail. This way, events can be specified so that the next multimedia data to be presented depend on the user's manifested interests. Applications like narrated virtual guided tours of urban landmarks (Fig. 2) can gain a lot from these capabilities, since the system can progressively provide more detailed information as the user naturally shows more interest in certain areas. Other applications of the Visorama System include interactive story telling in theme parks, virtual visits to real estate sites, and history education, among others.

More information is available on the Visorama Project Web page 2.

References
1 Binocular Omni-Orientation Monitor (BOOM). www.fakespace.com/product.html
2 The Visorama Project. www.visgraf.impa.br/visorama/s98

Figure1 The Visorama observation device. An artistic rendering (a). A photograph of the first prototype (b).

Figure 2
The Visorama prototype installed in a tourist site. The user is looking at a panorama of the surrounding area.

Remote Reality via Omni-Directional Imaging

Terrance E. Boult
EECS Department
Lehigh Univeristy
19 Memorial Drive West
Bethlehem, Pennsylvania 18015 USA
tboult@eecs.lehigh.edu
www.eecs.lehigh.edu/~tboult

Remote Reality (RR) is an approach to providing an immersive environment via omni-directional imaging. It can use live or pre-recorded video from a remote location. While less interactive than traditional VR, remote reality has important advantages: there is no need for "model building," and the objects, textures, and motions are not graphical approximations.

Omni-Directional Imaging

The development of the RR system was made possible by recent research,[1] which revolutionized wide-field-of-view imaging by introducing the paracamera, a system that directly captures a full hemisphere (or more) while maintaining a single perspective viewpoint. Because it captures the viewing hemisphere (or more) simultaneously, it can be used for full motion video. Furthermore, placing two paracamera systems back to back allows a true viewing sphere (360 x 360 viewing). Unlike fish-eye lenses, each image in the paracamera system can be processed to generate GEOMETRICALLY CORRECT perspective images in any direction within the viewing hemisphere.

The omni-directional imager combines an orthographic lens and a parabolic mirror, where the axis of the parabolic mirror is parallel to the optic axis. Because the lens is orthographic, entering rays are parallel. By definition, rays parallel to the axis reflect off a parabolic surface at an angle such that they virtually intersect at the focus of the parabolic surface. Thus the focus of the paracamera provides a single "virtual" viewpoint, which is enabling many novel applications, (see www.eecs.lehigh.edu/research/omni). The single virtual viewpoint allows for consistent interpretation of the world in any viewing direction. To generate a proper perspective image from the para-image, consider an "imaging array" in the desired viewing direction. For each pixel, logically cast rays through the focus and intersect the measured image. The resulting spatially varying resampling can be very efficiently implemented using spatial lookup tables. As the HMD turns or one "zooms in," the virtual viewpoint is stationary; only the virtual "imaging array" is moved.

The Remote Reality System

The main components of the system are the omni-directional camera, video recording systems, car mounting bracket, and a head-mounted-display (HMD). See www.cyclovision.com for commercial Paracameras with basic WindowsNT software. The stock paracamera is self-contained. Smaller custom designs for camcorders can be produced.

Because omni-directional imaging compresses a viewing hemisphere into a small image, maintaining resolution and captured image quality is quite important. While the process scales to any size imager, the current system uses NTSC (640x480) or PAL (756x568) cameras. Note that the "spatial resolution" of the image is not uniform. While it may seem counter-intuitive, the spatial resolution of the omni-directional images is GREATEST along the horizon,. The spatial resolution along the horizon is 4.2 pixels/degrees (5.1 for PAL).

The prototype system balances cost and quality. Our current data collection system was approximately $5,000, and the computing/HMD play-back system was about $3,000. The system uses a 233MMX CPU (running Linux) and video capture card. It computes monocular CIF-resolution, full-rate "video" (320 x2 40, 30 fps NTSC), which is reasonably well matched to the Virtual I-O glasses. The built-in head tracker provides yaw, pitch, and roll, with updates to the viewing direction at 15-30 fps. (A mouse or joystick can also be used for view selection.) We are currently adding GPS localization to the collection system to better support augmented reality applications.

Reference

1 S.K. Nayar. Catadioptric omnidirectional video camera. Proc. of DARPA Image Understanding Workshop, May 1997. (See also IEEE CVPR97)

Assisted Articulation of Closed Polygonal Models

Marek Teichmann and Seth Teller
MIT Laboratory for Computer Science
Room NE43-242
545 Technology Square
Cambridge, Massachusetts 02139, USA
marekt@lcs.mit.edu

We consider the problem of articulating models represented by polygonal meshes. Commonly, an articulated skeleton (also called an "I-K skeleton") is created, then bound to the mesh model by the user, typically by manual indication of a correspondence between elements of each structure. This process can, however, be tedious. We propose a system to automate both tasks.

The input to our system is a polygonal mesh of genus zero representing an object. We generate a tree-like skeleton with very little user intervention. The mesh is then automatically attached to the skeleton via a spring network. Subsequent user or keyframe-based manipulation of the skeleton results automatically in corresponding articulation of the original model.

The advantages of our system over previous approaches are that a skeleton is generated with little user input, as opposed to systems that require bone and muscle models,[2] or those that require an underlying lattice[3] or skeleton.[1] Furthermore, an easy extension to the system permits local variable rigidity of the skin (ex. for modeling scales or calluses), which would not be easy to simulate in classical skeleton-based deformation systems such as 1. On the other hand, our system currently lacks the ability to simulate the "muscle-bulge" of that system.

Medial Axis Approximation

First, the medial axis of the object is approximated as follows. The Voronoi diagram of the mesh vertices is computed, and the resulting skeleton is viewed as a graph. The user then selects a few points on the Voronoi skeleton to identify object features to be animated (e.g., at finger tips), these will be the leafs of the skeleton. The algorithm proceeds by identifying those vertices that can be removed without disconnecting the graph. These are obtained by computing the biconnected components of the current graph. Among those, only vertices of low importance (unselected vertices close to the mesh surface) are removed. This process is repeated until a tree is obtained, whose nodes will be anchor points for the mesh.

The resulting tree lies near the "center" of the model, but is connected to the user- selected tips. The resulting graph can be seen in the third image.

I-K Skeleton Construction

The user can then interactively add or remove points of articulation, as well as define local coordinate systems at each articulation. Chains of Voronoi skeleton vertices between articulation points become anchor points for the mesh, and each forms an I-K skeleton branch. The system then connects these articulation points to animation variables and provides a means to couple external animations (e.g., keyframe data) to animate the model via the skeleton. The resulting skeleton can be seen in the figure as yellow-red edges.

Mesh Animation

To connect the skin to the skeleton, we create a spring network that allows the skin of the model to follow deformations of the skeleton without breaks, folds, or excessive stretching in the mesh for reasonable motions of the skeleton. Each mesh vertex is connected to the closest anchor point on the skeleton. This is accomplished by computing the Delaunay triangulation of the mesh vertices and skeleton anchor points. An example for a simple leg model is shown in the last image.

At animation time, for each animation frame, the spring network is simulated until it stabilizes. The resulting mesh can be output for rendering. An example of a model and a new pose can be seen in the figure. We would like to thank Viewpoint Data Labs for providing models for this research.

References
1 Alias | Wavefront Reference Manual. Alias Wavefront 1996.
2 R. M. Koch, M. H. Gross, F. R. Carls, D. F. von Bueren, G. Fankhauser, and Y. Parish. Simulating facial surgery using finite element methods. SIGGRAPH 96 Conference Proceedings, 421-428. August 1996.
3 Ron MacCracken and Kenneth I. Joy. Free-form deformations with lattices of arbitrary topology. SIGGRAPH 96 Conference Proceedings, 81-188. August 1996.

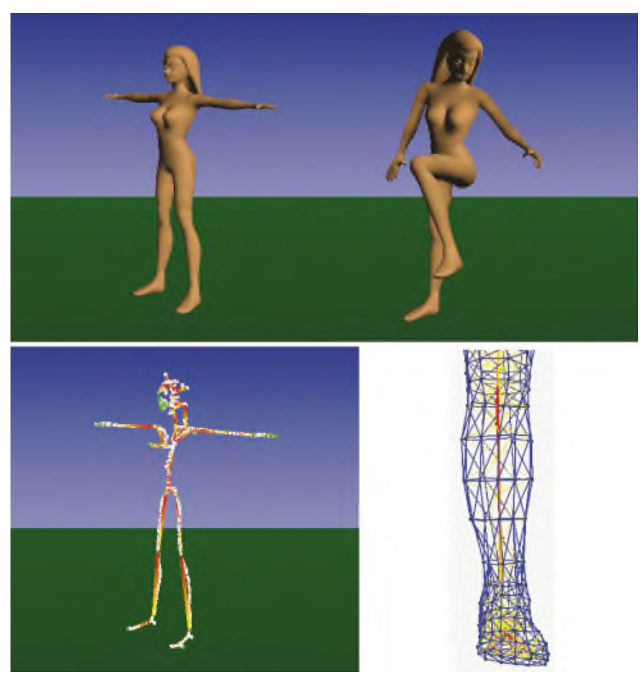

A model, a pose obtained during animation, and its skeleton. The last image shows the springs (yellow) attaching the skeleton points to the mesh (blue) of another, simpler model.

Dynamic Modeling of Human Hair and GUI Based Hair Style Designing System

Keisuke Kishi and Shigeo Morishima
Faculty of Engineering
Seikei University
3-3-1 Kichijoji-kitamachi Musashino
Tokyo 180-0001 Japan
keisuke@ee.seikei.ac.jp, shigeo@ee.seikei.ac.jp

Generation of super-realistic faces has focused recently on computer-human interaction or virtual people synthesis. Though the naturalness of hair is a very important factor in the visual impression, it is treated with a very simple model or as a texture. Because real hair has huge pieces and complex features, it is one of the hardest objects to model with computer graphics.

In this sketch, a new hair style modeling system is presented. The system helps designers use computer graphics to easily create any hair style with a tuft editor. Each piece of hair is modeled independently, and dynamic motion can be easily simulated by solving motion equations. However, each piece of hair has a few segments, which modeled with a 3D B-spline curve, so the calculation cost is not so large.

It is necessary to create a model of each piece of hair to generate the natural dynamic motion of hair blowing in the wind. However, flowing hair is very complex, and the average human head has more than 100,000 strands of hair. So storing a complete head of hair requires a huge polygon model.[1] In this system, each strand has a few segments to be controlled and is expressed with a 3D B-Spline curve.

In order to manipulate hair to produce any hair style, some of the hair pieces are treated as one tuft simultaneously. Each tuft is composed of more than seven squares, through which hair segments are passing. These squares are the control point of each hair segment, so manipulation of each square can be rotated, shifted, and modified to realize any hair style.

Each hair piece is modelled with a 3D B-Spline curve and as a very thin circle pipe.[2] So surface normal can be determined in each pixel to introduce the Lambert Phong models. To control hair thinness, alpha-blending methods are introduced. Blending hair into the background scene can make the hair sharper at the end.

Dynamic simulation of hair blowing in the wind is generated by calculation of the position of shape control points with external force. Each strand of hair is assumed to be composed of a few rigid stick segments, and a motion equation is solved from the root of each piece.

Our hair designing system makes it easy to manipulate complex hair styles by tuft modeling, and the GUI makes the designing process easier. Dynamic modeling of each strand of hair can generate realistic animation of hair blowing in the wind. In our future work, we will create straggling hair on the nape and hairlines, by appending new functions to the current system. We will also investigate collision detection among strands of hair .

References
1 Y. Watanabe and Y.Suenaga. A trigonal prism-based method for hair image generation. IEEE Computer Graphics & Applications, 47-53, January 1992.
2 S. Kobayashi, S. Morishima and H. Harashima. Representation of feel and motion of the thread-like objects. Proceeding of NICOGRAPH, No.6, 29-36, November 1990.

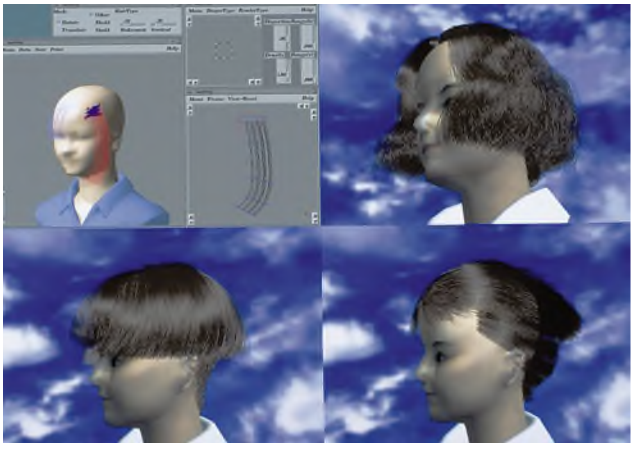

Session: Adventures in Moving

Directing Physics

Stephen Chenney
EECS Computer Science
University of California at Berkeley
387 Soda Hall, Number 1776
Berkeley, California 94720-1776 USA
schenney@cs.berkeley.edu

Session: Adventures in Moving

Physical models are an important source of simulations for training and animations but are notoriously difficult to direct. A training simulation may require that all objects are consistent with physics in their behavior and specify that a given object fall and land in a particular way. In typical current models these requirements are mutually exclusive, because initial conditions, surface geometry, and surface properties are exactly specified. Even if initial conditions are free, it is still a difficult problem to find those that lead to the desired outcome, particularly in collision-intensive systems, where the final state is extremely sensitive to initial conditions (real-world dice are used as random number generators for exactly this reason). For the same reason, most spacetime-constraint algorithms[1] cannot be applied.

Physically based simulations tend to lack the fine detail present in the real-world motion of objects, in part due to the simplifications inherent in modeling the world for simulation.[2] In our work, we enhance the physical model: normal vectors, restitution, and friction now vary as random fields over each surface. The random values used in the simulation make the overall animation random, and this randomness enhances visual plausibility. However, the animations can be made consistent with the physical world by ensuring that the random fields used are consistent with the properties of the objects being modeled, with the advantage that rough wooden tables will now behave "more randomly" than smooth glass tables.

This model corresponds more accurately to the physical world, and it can also be directed. We have one animation for each possible instantiation of the random surface values; to direct the system, we restrict our attention to those animations that satisfy a director's constraints. The probability of an animation is simply the product of the probabilities of all the physical parameters used, according to the random fields. To enforce constraints, we weight this probability according to how well the resulting animation satisfies the constraints. Any type of constraint is possible, provided it can be expressed as a positive weight that increases as the animation comes closer to satisfying the constraints. In our examples, we can constrain one or more of time, position, velocity, orientation, angular velocity, and the number thrown by a roll of the die.

To sample animations, we use a Markov chain Monte Carlo (MCMC) algorithm.[3] This generates an infinite number of animations, with the physical properties and resulting trajectory distributed according to the weighted probability described above. Hence, each animation is plausible according to our world model and satisfies the direction constraints.

The accompanying image shows a sample animation of a set of dice generated using the extended world model and the MCMC algorithm. Each die is shown where it came to rest as a result of a constrained physical simulation, with the entire trajectory of the last die indicated with a time-lapse sequence. At each collision, a surface normal vector, friction, and restitution coefficient were sampled from a spline surface, where the control points of the spline are random variables. To generate a new animation, we randomly perturb the control points of the splines, with a preference for those controls that influenced earlier parts of the animation. We also randomly perturb the initial conditions by small amounts. These six paths took approximately two hours of off-line time to compute, by starting from random animations (which did not satisfy the constraints) and iteratively generating new animations until the constraints were satisfied.

This approach is simple, general, and effective. It is particularly attractive because it generates many animations, all of which are close to satisfying the constraints and consistent with the physics of the world – an important consideration for training simulations and computer games.

References
1 Zicheng Liu, Steven J. Gortler and Michael F. Cohen. Hierarchical spacetime control. Computer Graphics: Proceedings of SIGGRAPH 94, 35-42, 1994.
2 Ronan Barzel, John F. Hughes, and Daniel N. Wood. Plausible motion simulation for computer graphics animation. Computer Animation and Simulation '96, Proceedings of the Eurographics Workshop in Poitiers, France, 184-197, 31 August 31-1 September 1996
3 Walter R Gilks, Sylvia Richardson, and David J Spiegelhalter. Markov chain Monte Carlo in practice. Chapman & Hall, 1996

Faster Integration of the Equations of Motion

Breton M. Saunders
Computer Laboratory
University of Cambridge
Cambridge CB2 3QG United Kingdom
bms20@cl.cam.ac.uk
www.cl.cam.ac.uk/users/bms20

Some of the most realistic motion in computer animation has been generated by physical simulation. The motion is generated by solving a system of differential equations approximately at each time step, and numerical methods are employed for this purpose. In practice, evaluation of the differential equations can be very time consuming, especially if physical constraints on the system are required. Common algorithms, such as Runge-Kutta, achieve their stability by approximating higher-order derivatives through multiple evaluations of the differential equations over a single time-step.[1] However, for computer animation we can usually tolerate a "visually acceptable" sacrifice in accuracy for faster evaluation of the equations of motion. In order to achieve interactive simulation rates, it is necessary to find explicit methods that meet the accuracy requirements and maintain stability while only requiring a single evaluation of the forces acting on the system per time step.

Interest in symplectic integration schemes has grown in recent years in the molecular dynamics community for the speed and accuracy they offer. One offshoot of symplectic integration schemes are leapfrog style integrators. These integrators require only one evaluation of the forces acting between particles (or rigid bodies) per time step and achieve their stability through formulation of the integrator rather than through multiple derivative evaluations.

Leapfrog integrators are derived from Hamilton's equations of motion and are used to integrate momentum and position separately. The Verlet method is an example of a symplectic leapfrog integrator that provides second-order accuracy and good stability for separable Hamiltonian systems.[2] However, with many cases in computer animation the system is neither Hamiltonian nor separable. In this sketch, I present the leapfrog methods and the extensionsto make them usable for computer animation work. I compare these methods with commonly used schemes and show how they can be used for real-time work with particle systems or rigid body simulation.

In each of the figures above a leapfrog method was employed to simulate the equations of motion at a fixed time-step of 1/30th of a second. Fig. 1 shows a particle constrained to lie on a ring. Initially, the particle was slightly off vertical and at rest. In comparison with the midpoint method, the modified Verlet scheme performed significantly better. Although both methods demonstrate energy gain, with the midpoint method the gain is so drastic that the particle spins over the top on the first revolution.

Fig. 2 shows a triple pendulum modelled as implicitly constrained particles using Witkin's method.[3] The modified Verlet scheme demonstrated good behaviour at one force computation per time step.

Fig. 3 shows a Hamiltonian system where a rigid body is attached to a spring. The equations of motion are simulated using Dullweber's method.[4] This is an example of a leapfrog method applied to a rigid-body system. This particular system is symplectic, and energy is conserved within a narrow band.

References
1 Press, Teukolsky, Vetterling and Flannery. Numerical recipes in C. Cambridge University Press, 1988.
2 Skeel, Zhang and Schlick. A family of symplectic integrators: Stability, accuracy, and molecular dynamics applications. SIAM J. Sci. Comput., Vol. 18, No.1, 203-222, Jan 1997.
3 Witkin and Baraff. Physically based modeling: Principles and practice. SIGGRAPH Course Notes #19, 1997.
4 Dullweber, Leimkuhler, and McLachlan. Symplectic splitting methods for rigid body molecular dynamics. Journal of Chemical Physics, #107, October 1997

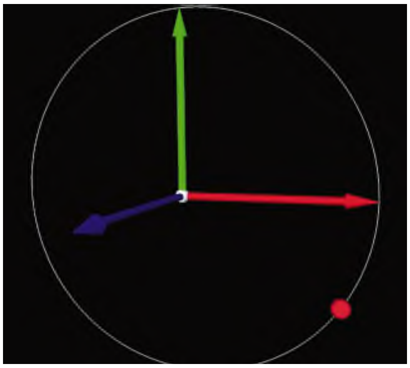

Figure 1

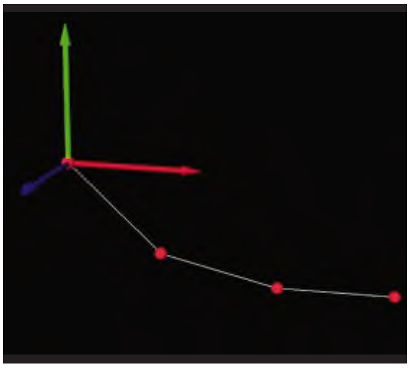

Figure 2

Figure 3

Physics-Model-Based 3D Facial Image Reconstruction from Frontal Images Using Optical Flow

Shigeo Morishima, Takahiro Ishikawa, Demetri Terzopoulos
Seikei University
3-3-1 Kichijoji-kitamachi
Musashino
Tokyo 180 Japan
shigeo@ee.seikei.ac.jp
www.ee.seikei.ac.jp/user/shigeo/

Physics-based face image synthesis is one of the most realistic approaches to realize life-like agents in computers. A facial-muscle model[1] is composed of facial tissue elements and muscles. In this model, forces affecting facial tissue elements are calculated by contraction of each muscle, so the combination of each muscle parameter determines a specific facial expression. Then each muscle parameter is specified in a trial-and-error procedure comparing the sample photograph and generated image using our Muscle-Editor to generate a specific face image. In this sketch, we propose a strategy for automatic estimation of facial muscle parameters from optical-flow using a neural network. This corresponds to 3D facial-motion tracking from a 2D image under the physics-model-based constraint. This technique is also 3D-motion estimation from 2D tracking in a captured image under the constraint of the physics-based face model.

We use optical flow of the facial image to measure transformation of the face when an expression appears. However, all of the flows are not used independently. They are compressed in the window depending on the position of each muscle. Optical flow is calculated by a block-matching method, and the block size is 8 by 8 pixels in an image size of 720 by 486 pixels. Then we can determine the difference between any specific expression and a neutral with 96 dimensional vectors including 48 windows with x-y directions.

The learning pattern is composed of a data pair: a muscle-parameter vector for the output layer and an optical-flow vector for the input layer. Neural networks were trained using back propagation. They include 55 images (i.e., five keyframes for transition from neutral to each of six basic expressions and mouth shapes of five vowels. Their learning pattern is not given at once but gradually increased to escape the local minimum.

To confirm whether the learning process of the neural network is successfully completed, we input the learning data into the input layer of the neural network and resynthesize the facial image using muscle parameters from the output layer of the neural network. Example results are shown in the figures. There are slight differences in fine detail between the original and re-generated images, but subjectively, the overall facial features and expressions are almost the same as the original image. So the mapping rules work well for the training data. Muscle parameters are decided only from the 2D image, but the 3D facial image is well regenerated.

Optical flows in image sequence are given to the neural network frame by frame, and a facial animation is generated by the muscle parameter sequence from the output layer. This tests the interpolation effect on our parameter-conversion method based on the generalization effect of the neural network. Training data include only five keyframes in the transition from neutral to each of six basic expressions. After our subjective evaluation, overall facial features are continuously generated between the keyframes to make the motion natural, and the impression is very close to the original.

Currently, our facial-muscle model requires significant computation time. As a result, this method is not real-time processing. Furthermore, from our evaluation we can see that there are limits to generating arbitrary expressions with our model, so relocation of the muscles and a new definition of the physics for the new muscles is the next subject of our research. The most important future problem is how to evaluate the resynthesized face expression quantitatively.

Reference
1 Yuencheng Lee, Demetri Terzopoulos, and Keith Waters. Realistic modeling for facial animation. Proceedings of SIGGRAPH 95, 55-62, 1995.

Vector Field Comparisons using Earth Mover's Distance

Yingmei Lavin, Rajesh Batra, Lambertus Hesselink
Stanford University
143A Escondido Village
Stanford, California 94305 USA
yingmei@kaos.stanford.edu
raj@kaos.stanford.edu
bert@kaos.stanford.edu

Vector field topology has received much attention in both the visualization and the physical science communities. However, to our knowledge, almost no work has been done on quantitative measurements for vector-field comparisons. In this sketch, a novel approach is introduced to define a quantitative measure of closeness between vector fields. The usefulness of this measurement can be seen when comparing computational and experimental flow fields under the same conditions. Furthermore, its applicability can be extended to more cumbersome tasks such as navigating through a large database searching for similar topology.

This new measure relies on the use of critical points, which are key features in vector-field topology. In order to characterize critical points, α and β parameters are introduced. In Clifford algebra, a linear vector field can be expressed as $\vec{v} = (az + b\bar{z} + c)e_1$ where $a = a_1 + a_2 i$, $b = b_1 + b_2 i$ and $z = x + yi$.[1] Eigenvalues of the Jacobian around its critical point are $\lambda_{1,2} = b_1 \pm \sqrt{|a|^2 - b_2^2}$. And $\alpha = b_1$, $\beta = sign(|a|^2 - b_2^2)\sqrt{||a|^2 - b_2^2|}$ are then used to form a closed set of eight unique patterns for simple critical points, as shown in Figure 1. These patterns are also basic building blocks for higher-order nonlinear vector fields. A new α-β space is introduced where these basic patterns are points located at particular values of (α,β). Vectors in this space obey all the rules defined for a regular *2D* Euclidean space. The Earth Mover's Distance(EMD), first introduced by Yossi Rubner et al. for fast retrieval of similar images in a large data base,[2] computes the minimal amount of work that must be performed to transform one distribution into the other by moving "distribution mass" around. The feature distribution, in our case, is defined as a set of α,β parameters for every critical point in a vector field. And a measure of distance, defined as $d_{ij} = \sqrt{(\alpha_i - \alpha_j)^2 + (\beta_i^2 - \beta_j)^2}$, is introduced to calculate the difference between a given set of vector fields. In order to get unbiased distance, α's and β's are normalized. This normalization process maps the basic patterns of critical points onto a unit circle in α - β space. After extracting the feature distribution, a set of EMDs are computed. Multidimensional scaling (MDS) was originally used for color perception and spatial frequency discrimination. Here, with the help of MDS, we can display the similarities and differences between selected vector fields in Cartesian space. This approach quantitatively measures the similarity and dissimilarity between vector fields. It is ideal for data compression of a large-flow field, since only the number and types of critical points along with their corresponding α and β parameters are necessary to reconstruct the whole field. It can also be used to better quantify the changes in time-varying data sets.

As an example, a group of 16 vector fields (figure 3) will be compared. The fields with similar distributations are clustered and those that differ will be distiguished spatially.The first four vector fields contain saddle points in various configurations. Therefore, it is expected that the EMD will cluster these topologies together. Vector fields 5, 6, 7, and 8 contain various topological elements such as repelling and attracting foci and nodes. Vector fields 9 and 13 both have a pair of foci and one center, but slightly different α and β values. Therefore, the distance between them should be very small. Vector fields 10, 11, and 12 all have one center and two nodes, so they are expected to group together. Vector fields 14, 15, and 16 all have two foci but different α and β values, particularly for vector field 15, where the α value is so close to zero that two foci appear like two centers. Results are shown in Figure 4, and, as can be seen, they agree well

with our expectation: the $\binom{16}{2}$ EMDs were computed, and

MDS correctly grouped the similar vector fields together and separated the dissimilar ones according to the distances between them. (The vector fields in Figure 4} have been slightly moved for viewing purposes, because the distances between some vector fields are so small that they lay on top of one another.)

References
1 Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. A metric for distributions with applications to image databases. Proc. IEEE International Conference on Computer Vision,1998.
2 G. Sheuermann, H. Krüger, M. Menzel, and A. Rockwood. Visualization of higher order singularities in vector fields. Proc. IEEE Visualization, 67-74, 1997.

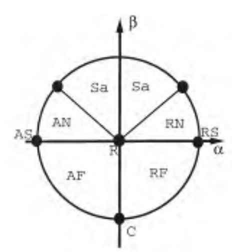

Figure 1 Basic patterns for critical point in (formula) space; *C* for center, *RN* for node, *AN* for attracting node, *RF* for repelling focus, *AF* for attracting focus, *St* for star, *Sa* for Saddle and *R* for fregular point

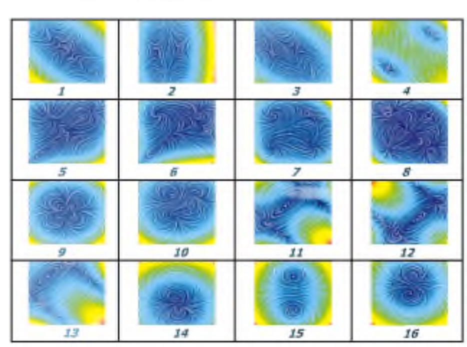

Figure 3 Sixteen different vector fields

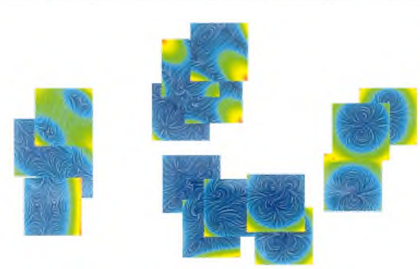

Figure 4 Sixteen different vector fields forted by MDS

Measuring Volumetric Coherence

Yuriko Takeshima
Ochanomizu University
Tokyo 112-0012 Japan
yuriko@imv.is.ocha.ac.jp

Session: Creating Coherence

An estimate of the degree of volumetric coherence can be obtained by examining how often the possible pairs of field values in a given regular, volumetric data set occur in the set of relative positions.

Let $\delta \equiv (\Delta x, \Delta y, \Delta z)$ be a 3D displacement, and let P_δ be a co-occurrence matrix whose (i, j) element represents the frequency that a voxel (x_0, y_0, z_0) having field value i occurs in position delta relative to a voxel (x_1, y_1, z_1) having field value j ($f_{min} \leq i, j \leq f_{max}$). Then, let P be the average of P_δ's for the set of displacement deltas of a given size τ in various directions. Here, 26 neighbors around a voxel (the Hamming distance $\tau = 1$) are examined for P to obtain isotropic information about the spatial distribution of field values in a volumetric data set.

As a candidate for volumetric coherence measure, the following statistic is derived from the averaged co-occurrence matrix P:

$$VCM \equiv -\sum_{i=0}^{n_l-1} \sum_{j=0}^{n_l-1} \frac{(i-j)^2}{{n_l}^2} P(i,j) \log_{10} P(i,j),$$

where n_l denotes the number of possible integer field values.

The statistic VCM can indicate the degree of volumetric coherence from the viewpoint of gray-scale uniformity of field values, since the two types of the spatial distribution expressiveness, Haralick's 2D statistics for textures, that is, contrast and entropy,[1] are retained. Obviously, the more coherent a given data set becomes, the smaller the VCM value becomes.

Feasibility of VCM

In order to demonstrate the feasibility of VCM as the volumetric coherence measure, preliminary experiments were performed with an H_2 electron density dataset and six modified versions of the Chapel Hill testbed volumetric data sets.[2] All of the datasets were normalized onto one byte unsigned integers (0--255). The platform used in the experiments was an SGI O2 system (CPU: R5000, Clock: 180MHz, RAM: 64MB).

Fig.1 illustrates to what degree the measure VCM correlates with several isosurfacing-related statistics in terms of complexity and accuracy. All of the statistics were derived by averaging over 2^8 isosurfaces with integer target values. The monotonicity of the curves shows that the VCM correlates strongly to:

- Percentage of cubes intersecting isosurfaces (cube hit rate).
- The number of isosurface patches per cube.
- Percentage of ambiguous cube sides.
- Percentage of topologically correct isosurface patchs with disambiguation schemes.

This result can be used as a guideline to resolve the temporal complexity vs. topological consistency trade-offs among the auxiliary schemes that attempt to alleviate the topological ambiguities of isosurfaces extracted by the Marching Cubes algorithm.[3] The authors are currently attempting to develop a VCM-based time-critical environment for indirect volume visualization.

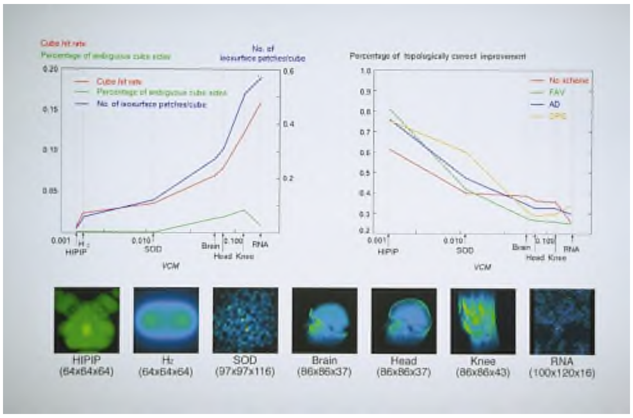

Figure 1 Correlation of the measure VCM to isosurfacing statistics: (a) spatial complexities and (b) accuracy of disambiguation schemes with respect to quadratic-fit gradient. (FAV: Facial Average Values, AD: Asymptotic Decider, CPG: Center-Pointing Gradient)[4]

References

1 Azriel Rosenfeld and Avinash C. Kak. Digital Picture Processing (2nd edition), volume II, chapter 12, 295--304. Academic Press, 1982.
2 Louis Noodleman and David Case (University of North Carolina, Chapel Hill). (ftp://ftp.cs.unc.edu/pub/projects/softlab/CHVRTD/voll).
3 William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. ACM Computer Graphics (SIGGRAPH 87 Conference Proceedings), 21(4):163-169, July 1987.
4 Allen Van Gelder and Jane Wilhelms. Topological considerations in isosur face generation. ACM Transactions on Graphics, 13(4):337-375, October 1994.

Artificial Evolution of Implicit Surfaces

Edward J. Bedwell and David S. Ebert
Computer Science and Electrical Engineering Department
University of Maryland Baltimore County
1000 Hilltop Circle
Baltimore, Maryland 21250 USA
ebedwe@cs.umbc.edu

Procedural techniques allow algorithmic and mathematical methods to generate complexity to a nearly infinite level of detail. By giving the user parametric control of procedural methods, extremely complex models can be created in a fraction of the time it would require to build them by hand. This methodology has become widely accepted throughout the computer graphics community, as evidenced by the work of Lindenmayer and Prusinkiewicz in procedurally simulated plant growth,[5] and the vast work that has been done in the realm of implicit surfaces.[1,3]

However, procedural techniques introduce new levels of difficulty. Creating these procedural tools often requires a strong understanding of complicated mathematics. Furthermore, even if one does have the mathematical knowledge necessary, providing parametric controls to one who does not is a non-trivial task. Sims presented genetic programming as a viable solution to this problem with his procedural texture generation system.[6] Genetic programming has also been proven useful to computer graphics. Dawkins made use of the genetic metaphor to generate 2D branching structures.[2] Gritz and Hahn applied genetic programming techniques to generate various procedural animation methods.[3] And Todd and Latham have combined genetic programming with solid geometry to produce amazing evolutionary art.[7] This sketch presents a system that combines implicit surfaces, as modeling primitives, with genetic programming to facilitate automated generation of exceedingly complex models. The class of implicit surfaces this system utilizes is characterized by a function $f : R^3 \rightarrow R$ that assigns some value to each point in three-space. The surface is the set of points such that $f(x,y,z) = 0.4$ A few predefined surfaces form our primitive community. The initial parents, from whom the system will generate offspring, are selected from this community. From each new generation the user can again select two individuals to mate, creating the next generation. This process is repeated until a surface is generated that fits the user's needs.

Implementation

To mate two surfaces, the user begins by selecting two parent individuals from the primitive community that exhibit traits wanted by the user. The functional representations are manipulated as LISP expressions. A series of LISP routines generates hierarchical parse trees from each expression. Mating is facilitated by a "cross-over" procedure in which subtrees are pseudo-randomly selected from each parent and swapped. Thus, two children are "born" to the user-selected parents. This process is repeated until the generation reaches a suitable size.

The newly generated expressions are converted to infix notation and embedded in ASCII scene description files. These scenes are then rendered using a modified version of POV-Ray, a publicly available ray-tracing package.

Results and Application

In biological systems, offspring created through the mating process will exhibit traits from both parents. This holds true in the artificial mating of implicit surfaces. The system has shown that even several generations will maintain traits from the original parents. Given a sufficiently large and diverse population of surfaces from which to begin, users should be able to create a surface that meets their needs.

Conclusion & Future Work

This system successfully demonstrates that genetic programming greatly simplifies the generation of complex implicit surfaces. Furthermore, the genetic metaphor allows for the transmission of desirable traits from each parent to the child. Thus, the user maintains some control of the evolutionary process. Future work includes an extended primitive set, differing probabilities for crossover points in the parse trees, mutation during the mating process, and blending multiple implicit surfaces through genetically determined blending functions.

Acknowledgments
Ryoichi Suzuki, of the Electrotechnical Laboratory in Japan, for providing the isosurface modifications to the Persistence of Vision raytracer.
Andrew Guy and Brian Wyvill of the University of Calgary, for making their software available.

References
1 J. Bloomenthal, C. Bajaj, J. Blinn, M. Cani-Gascuel, A. Rockwood, B. Wyvill, G. Wyvill. Introduction to implicit surfaces. Morgan-Kaufmann, 1997.
2 R. Dawkins, The blind watchmaker. Harlow Logman 1986.
3 L. Gritz, J. K. Hahn. Genetic programming for articulated figure motion. The Journal of Visualization and Computer Animation, Vol.6, 1995, 129-142.
4 J. C. Hart. Ray tracing implicit surfaces. WSU Technical Report EECS-93-014, 1993.
5 P. Prusinkiewicz, A. Lindenmayer, and J. Hanan. Developmental models of herbaceous plants for computer imagery purposes. Computer Graphics, Vol.22 No.4, 1988, 141-150.
6 K. Sims Artificial evolution for computer graphics. ACM Computer Graphics, 25(4), July 1991, 319-328.
7 S. Todd, W. Latham. Evolutionary art and computers. Academic Press, 1992.

A Progressive Global Illumination Solution Considering Perceptual Factors

Vladimir Volevich, Andrei Khodulev, and Edward A. Kopylov
Russian Academy of Sciences

Jerzy Sas
Technical University of Wroclaw

Karol Myszkowski
The University of Aizu
k-myszk@u-aizu.ac.jp
www.u-aizu.ac.jp/~k-myszk/progress

Realistic image synthesis usually involves some iterative methods for solving the global illumination problem. The fully converged solution may be quite time consuming for complex scenes, but taking into account basic properties of human perception such as poor sensitivity to the absolute luminance values, high quality images can be obtained on the order of single minutes or seconds using physically-based partial solutions. In this sketch, we describe a progressive technique designed along these guidelines.

We use a hybrid of stochastic and deterministic techniques. At first, a Monte Carlo Photon Tracing algorithm is used in which photons are bucketed directly into a mesh (making immediate image display possible). To reduce the problems of low-frequency noise at early stages of computation [1], the adaptive filtering of illumination at the mesh vertices is performed, taking into account illumination in the local neighborhood. If we were to relay on the neighborhood relations between the mesh elements, the result would strongly depend on the geometrical model and would not provide complete information for separate objects. Instead, we build a static and balanced 3D kd-tree structure for all vertices. The filtering is used at the intermediate stages of computation, and filter size is adaptively reduced as variance of the illumination estimate decreases for a given vertex. Figure 1a shows the result of the particle tracing phase after 6 seconds of computations (Pentium-200 MHz) for the scene containing over 17,000 polygons.

The appearance of images obtained at the first stage of lighting simulation approximates well the final images with the exception of views that contain many areas with strong direct lighting. To overcome these drawbacks, deterministic calculations of direct lighting are performed. In contrast to traditional approaches, the adaptive mesh subdivision for the deterministic calculations of direct lighting can be based upon predicted visible differences using the available stochastically-derived estimates of indirect illumination. Here we introduce perceptually-based mesh splitting criteria. For each mesh vertex we transform the stimulus luminance values to predicted perceived brightness using Stevens' power law, and a decision on the mesh splitting is made based on the local differences in brightness. Figure 1b shows a result of this stage of lighting simulation which took about 30 seconds.

The next stage of computation replaces the finished direct computations with additional stochastic particle tracing, which is then performed until a criterion accuracy is reached (Figure 1c). The final images can be rendered using ray tracing.

Another way in which perceptual factors were considered in this research was to use the Visible Difference Predictor, or VDP [2], to generate quantitative measures of the differences in the appearance of two images. Figure 1d shows, as a function of computation time, the percentage of pixels in the two images for which the differences should be detected at high probability (refer to [1] for visualization of the VDP predicted differences). As can be seen, the current technique converges quite rapidly.

References:

1 http://www.u-aizu.ac.jp/~k-myszk/progress - the Web page that provides more information on this work.

2 Scott Daly. The Visible Differences Predictor: An algorithm for the assessment of image fidelity. In A.B. Watson, ed., Digital Image and Human Vision, pp. 179--206. MIT Press 1993

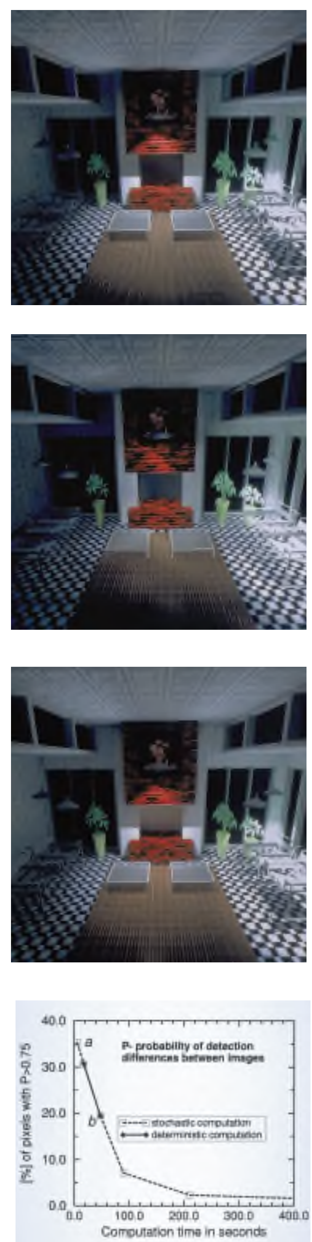

Figure 1 Progressive rendering:

a) Photon tracing (6 seconds)

b) Photon tracing (18 seconds) + deterministic direct lighting (30 seconds)

c) Fully converged solution (2 hours)

d) "Perceptual convergence" of the image quality as a function of computation time. The predicted differences between the intermediate images and that for the fully converged solution are computed using the VDP. Images a and b in this figure correspond to the time points *a* and *b* marked in the graph.

Visibility Driven Hierarchical Radiosity

Fredo Durand, George Drettakis, Claude Puech
iMAGIS/ GRAVIR-IMAG
BP 53
38041 Grenoble Cedex 09 France
[Fredo.Durand|George.Drettakis|Claude.Puech]@imag.fr
www-imagis.imag.fr/

Hierarchical radiosity (HR) simulations rapidly compute light exchanges in a given scene by subdividing the original objects to create an adaptive mesh representing light variations and shadows. Previous mesh subdivision strategies still suffer, however, from the limitations of regular (quadtree) subdivision. Discontinuity meshing solutions are usually limited to primary illumination and become intractable for more than two or three light sources. Visibility calculation is still done largely by ray-casting approximations, which can often produce erroneous results.

In this sketch, we present a novel visibility driven hierarchical radiosity algorithm by extending the visibility skeleton.[1] Our approach uses the extended skeleton to compute accurate visibility at vertices, to insert important discontinuities, and to effect an efficient refinement strategy. We use hierarchical triangulations and introduce a novel push-pull procedure for accurate display.

The most significant related work includes that of Lischinski et al,[5] which combines discontinuity meshing and HR. Discontinuity meshing is performed for a small number of direct lights, and a second pass is used to accurately compute visibility and illumination at vertices. Our visibility-driven HR algorithm is also related to wavelet radiosity[4] and to the work on perceptually based refinement.[3]

Extended Skeleton

The visibility skeleton is an encoding of all visibility events in a scene, using a graph structure. The nodes of the graph are extremal stabbing lines and the arcs critical line swaths that allow calculation of the exact visible contour of every polygon from every scene vertex.[1] We have extended this structure, thereby allowing calculation of views at vertices resulting from the subdivision of original scene polygons required by the radiosity algorithm. We have also replaced the quadratic table representation of arcs by linking mutually visible objects. These links are also used in the lighting solution. All polygon-vertex form factors (including those for indirect light transfers) are thus computed exactly.

Embedded Triangulation Hierarchies for Accurate Hierarchical Radiosity

To accurately represent radiosity across surfaces, especially for the very irregular discontinuities created by shadow borders, we have chosen to use hierarchical Delaunay triangulations (Similar to 2). Instead of computing radiosity at patch centers, we gather light at each vertex. This requires a new push procedure, which allows the contribution of a source to a vertex at a certain level to be pushed to the corresponding children. This is done by storing radiosity differences between the approximation at the higher level and the value at the current level. During the push, a vertex value at the current level is simply the sum of the stored difference and the interpolated value of the higher level at this vertex.

The result of this approach is a high-quality, continuous linear representation of radiosity at the leaves of the mesh, obviating the need for a second pass for display.

Visibility Driven Hierarchical Radiosity

We maintain two link types: vertex-polygon links for gathering and polygon-polygon links used to determine if a light transfer should be refined. The two link types are subdivided simultaneously, resulting in a consistent representation of light transfer.

To refine the scene polygons, we use the visibility information of the skeleton. For a given receiver, we examine all its polygon-polygon links and the contained visibility events. The corresponding discontinuities are ranked using their perceptually based importance, and only a certain percentage are actually inserted into the receiver Delaunay triangulation. It is important to note that this operation takes into account all visibility interactions, including those for indirect illumination links. We also subdivide at the maximum of the unoccluded light source radiosity for large receivers. Subdivision is based on the difference between exact and interpolated values for the unoccluded case and an estimation of shadow importance elsewhere.

Results

The following figure shows the result of our implementation for a scene containing 10 light sources and 486 polygons, and a scene of 534 polygons with mainly indirect lighting due to the lefthand table and the bed. The mesh adapts to the combined shadow effects automatically, resulting in high-quality illumination. The images were rendered directly using graphics hardware and did not require a local or final gather pass.

References

1. F. Durand, G. Drettakis, and C. Puech. The visibility skeleton: A powerful and efficient multi-purpose global visibility tool. In Computer Graphics Proceedings, Annual Conference Series, 1997, 89-100, 1997.
2. L. De Florani and E. Puppo. Hierarchical triangulation for multiresolution surface description. ACM Transactions on Graphics, 14(4):363-411, October 1995.
3. S. Gibson and R. J. Hubbold. Perceptually driven radiosity. Computer Graphics Forum, 16(2):129-141, 1997.
4. S. J. Gortler, P. Schroder, M. F. Cohen, and P. Hanrahan. Wavelet radiosity. In Computer Graphics Proceedings, Annual Conference Series, 221-230, 1993.
5. D. Lischinski, F. Tampieri, and D. P. Greenberg. Combining hierarchical radiosity and discontinuity meshing. In Computer Graphics Proceedings, Annual Conference Series, 199-208, 1993.

Figure 1 a) Skeleton construction took 2 min 23s.; meshing/lighting step took 8 min. (b) Skeleton construction took 4 min 12 s; meshing/lighting step took 7 min.

Session: The Light Brigade

Radiance Maps: An Image-Based Approach to Global Illumination

Philipp Slusallek, Wolfgang Heidrich, Hans-Peter Seidel
Universität Erlangen
Computer Graphics Group
Am Weichselgarten 9
D-91058 Erlangen, Germany
slusallek@informatik.uni-erlangen.de
www9.informatik.uni-erlangen.de/Persons/Slusallek

We present a new image-based approach to global illumination that is designed to explicitly make use of coherence in the radiance field of a scene. The approach builds on recently developed ideas from image-based rendering in order to reuse previously computed geometric and illumination information as far as possible.

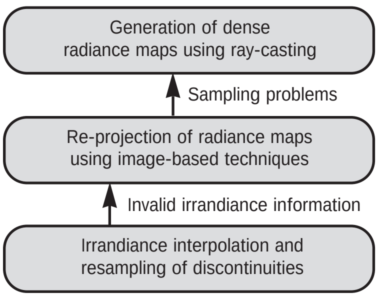

Figure 1
The overall structure of our approach. If insufficient information is available at one layer, the algorithm proceeds to the next one.

The algorithm consists of a hierarchy of three methods, each taking advantage of a different coherence pattern in the radiance field. At the top level, we sample primary layered-depth radiance maps of the environment, thereby abstracting from the exact geometry of the scene. In the intermediate level, coherence in ray space is exploited in order to derive new radiance maps by re-projection from primary maps. At the lowest level, gradient information is used to spatially interpolate smooth irradiance contributions, handling non-smooth contributions separately. Missing or uncertain information in lower levels is handled by resampling the environment and feeding these samples back to adjacent top-level maps. This adaptively improves the layered-depth images and thus the results of later re-projection operations.

The essence of our algorithm is an image-based approach to quickly provide accurate radiance maps anywhere in a scene. In this sketch, we use the radiance maps to compute irradiance in diffuse environments. However, the comprised information can be used for a much wider range of applications.

Overview of the Radiance Map Algorithm

As its main source of information, the algorithm uses radiance maps consisting of radiance and depth samples of the environment, relative to some point of interest, thereby abstracting from the exact geometric representation of the scene. The approach makes use of a combination of three different algorithms working together at different levels (see Fig. 1), each of which is designed to exploit a different aspect of coherence in the radiance field. Together, they form a global illumination algorithm that is fast, robust, and efficient both in terms of memory and time.

Lowest Level: Irradiance Interpolation

At the lowest level, the spatial coherence of irradiance (radiance integrated over all directions) is exploited by interpolating it from previously computed nearby irradiance samples using available gradient information. This part is based on previous work by Ward.[4] However, we use the information and the flexibility provided by the other two levels to interpolate only the smoothly varying irradiance contribution. Interpolating those contributions that would cause strong variations is explicitly avoided and handled separately using information from the upper levels of the algorithm.

Intermediate Level: Reprojection of Radiance Maps

If insufficient irradiance information is available (too few samples or too little coherence between them), we proceed to the intermediate level, where we directly reuse previously computed primary radiance maps. These maps represent views of the environment from particular locations. By reprojecting several of them towards the current position, we can generate a new view,[1,3] which we refer to as a secondary radiance map. Reprojection efficiently uses the inherent coherence in ray space and also improves the effective sampling density of the environment by combining several maps. In contrast to pure image-based methods, we can use the available scene description for resampling small amounts of missing information by ray casting where necessary.

Figure 2 Ward's office scene computed with indirect illumination only.

Finally, the algorithm at the top level is responsible for obtaining primary radiance maps. It initially samples the incident radiance and the distance to geometry on a regular directional grid and refines these samples in regions of high variance. We obtain samples not only from visible surfaces, but also from occluded geometry. These additional samples are added to primary radiance maps, following the idea of layered-depth images (LDIs).[2] LDIs are refined later by integrating the information obtained during resampling of missing information in the bottom two levels.

Fig. 2 shows the result of computing the indirect illumination with the new algorithm in a non-trivial scene.

References
1 Shenchang Eric Chen and Lance Williams. View interpolation for image synthesis. Computer Graphics (SIGGRAPH 93 Proceedings), 279-288, August 1993.
2 Steven J. Gortler, Li wei He, and Michael F. Cohen. Rendering layered depth images. MSR-Tech-Report MSTR-TR-97-09, Microsoft Research, March 1997. available from www.research.microsoft.com/research/graphics/cohen/.
3 Leonard McMillan and Gary Bishop. Plenoptic modeling: An image-based rendering system. Computer Graphics (SIGGRAPH 95 Proceedings), 39-46, August 1995.
4 Gregory J. Ward and Paul S. Heckbert. Irradiance gradients. In Third Eurographics Workshop on Rendering, 85-98, May 1992.

Interactive Modification of Real and Virtual Lights for Augmented Reality

Céline Loscos, George Drettakis and Luc Robert
iMAGIS/GRAVIR-IMAG
BP 53
38041 Grenoble Cedex 09, France
Celine.Loscos@imag.fr
www-imagis.imag.fr/

Computer-augmented reality (CAR) provides the ability to mix real and virtual worlds, and is used in entertainment, television, or to augment systems for repair and manufacturing. Reconstructed real-world models based on video images obviate the need for the often-arduous task of modeling complex scenes in detail. We focus here on common illumination, the interaction of shadows and other lighting effects between real and virtual objects. We build on the approach of Drettakis, Robert, and Bougnoux,[1] for which vision techniques are used to extract an approximate model of the scene. Based on work by Fournier et al.,[3] a basic representation of illumination is built using a radiosity algorithm. Virtual objects can be inserted, using the incremental hierarchical radiosity (HR) approach of George Drettakis, and Francois Sillion.[2]

In this sketch, we present novel solutions that allow addition of virtual lights and modification of real light source intensities. The main problem is that real-world lights cast shadows that are embedded in the textures of the model used to reconstruct the real world. No previous solutions exist to the problem of identification and removal of these shadows, thus permitting light modification.

Our solution is based on a pre-processing step that modifies the textures by weighting them according to shadow relationships identified in the radiosity solution. A correction is then applied to reinstate an approximation of the original shadows. Real lights can then be turned off or their intensities changed interactively, and virtual lights and objects can also be rapidly inserted.

Initialization of the Radiosity System

In [3] and [1], radiosity parameters are retrieved from the textures extracted from the real scene images. In particular, real light source emittance and object reflectance is estimated based on texture values and simple assumptions on average scene reflectance and radiosity. After this estimation, the geometry representing the real scene is subdivided into patches with radiosity B_i^o (original radiosity). After adding a virtual object, modified radiosities (i.e., in the shadow of a virtual object), are B_i. The texture t_i of each patch is modulated by the ratio

$$\frac{B_i}{B_i^o} \text{ (Eq. (1)).}$$

The ratio of Eq. (1) is insufficient to remove shadows when we change the real light intensities, since the texture t_i already contains the shadow due to the real sources. To perform this correctly, we need to remove the shadow from the texture and represent it explicitly using the radiosity solution.

For each patch in shadow (as determined by the radiosity system), we compute a value $\bar{E}_i$, which is the irradiance which would have arrived at patch i, but was blocked by occlusion. To generate a modified texture, we compute a correction factor based on $\bar{E}_i$ patches in shadow. Unoccluded patches are displayed using the ratio of Eq. (1), which has unit value, since no virtual object has yet been inserted. The resulting image will be used to replace the original texture.

For patch i and s = 1..N sources with radiosities B_s, we use the following correction factor:

$$\frac{\rho_j \Sigma_{s=1..N}(F_{si} B_s + \bar{E}_i)}{B_i^o} \text{ X } t_i$$

where B_i^o is the original radiosity value and F_{si} the form-factor from s to patch i. We scale by the reflectance ρ_j of a *user-selected* patch, which is not in shadow and has the appropriate color (of the shadowed surface).

For each patch corresponding to the modified texture, we update the reflectance as before and set the original radiosity to $B_i^o = \rho_i (E_i + \bar{E}_i)$, capturing the radiosity changes from the texture.

The radiosity is then recomputed by solving the HR system. To compensate for approximation errors, a display correction factor is used based on previous texture values for modified patches.

Interactive modification of Real Lights

We can now interactively modify the illumination of the real light sources. While maintaining links and subdivision, we first remove the radiosity and iteratively gather and push-pull, thus updating radiosity based on the new emittance value of a source. For example, if we switch off a light, the corresponding shadows are no longer perceptible in the image. We can also insert a virtual light source in the hierarchy and project the shadows of the real scene objects. Fig. 4 shows the initial illumination on the left, and on the right, the same real scene with all of the four real light sources switched off and a virtual light source switched on.

Conclusion and Future Work

We have developed a solution to remove real objects' shadows and perform interactive modifications of both real and virtual illumination. Shadows could be more accurately delimited during texture modification using a color-based refinement.

Our method could be extended to remove real objects and insert virtual replacements, since shadow removal is the hardest problem in this case. Since the reconstruction of the real-world model is never perfect, we expect to increase user interaction to dynamically determine the level of geometric detail required during an interactive session.

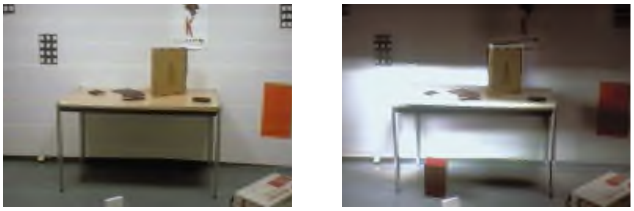

Figure 1 (a) Initial illumination, (b) modified illumination

References

1 George Drettakis, Luc Robert, and Sylvain Bougnoux. Interactive common illumination for computer augmented reality. In J. Dorsey and P. Slusallek, editors, Rendering Techniques '97, 45-56, 8th EG workshop on Rendering, Saint Etienne, France, June 1997. Springer Verlag.

2 George Drettakis, and François Sillion. Interactive update of global illumination using a line-space hierarchy. In Turner Whitted, editor, SIGGRAPH 97 Conference Proceedings, Annual Conference Series, 57-64. ACM SIGGRAPH, Addison Wesley, August 1997.

3 Alain Fournier, Atjeng S. Gunawan, and Chris Romanzin. Common illumination between real and computer generated scenes. In Proceedings of Graphics Interface '93, 254-262, Toronto, May 1993. Canadian Information Processing Society.

Fast Multi-Layer Fog

Justin Legakis
Computer Graphics Group
Laboratory for Computer Science
Massachusetts Institute of Technology
legakis@mit.edu
graphics.lcs.mit.edu/~legakis

Fog is an important element of real-time computer graphics systems. It is an inexpensive way to add realism to a scene, and it can provide important depth-cuing information. In addition, visual-simulation applications can take advantage of the occlusion provided by fog and cull objects that are too far away to be seen clearly.[3]

In this work, the standard fog model is extended from fog of uniform density to fog whose density varies as a function of height. The computation of the attenuation due to fog is still performed in constant time, independent of the complexity of the fog density function.

Fog rendering uses the attenuation projection[2] equation borrowed from volume rendering:[1]

$$fog = e^{-\int_a^b \alpha(t)dt} \tag{1}$$

The standard model used in real-time computer graphics systems, such as OpenGL, is fog of uniform density throughout space. For a given surface point, the integral in equation 1 is simply proportional to the distance from the eye to the point. This distance is usually approximated by z in eye space.[3]

Using a more general fog model requires integration of the fog-density function along the line from the eye to the surface. The first step of deriving the fast multi-layer fog algorithm takes advantage of the fact that the fog varies only as a function of height. The problem can be reduced to three variables: Y_1, the height of the camera; Y_2, the height of the surface point; and ΔD, the horizontal distance between the two points (points a and b in Figure 1):

$$\Delta D = \sqrt{(X_E - X_S)^2 + (Z_E - Z_S)^2} \tag{2}$$

Furthermore, the integral in equation 1 over the line ab can be written in terms of the integral over the line ac, which depends only on Y_1 and Y_2:

$$\int_a^b \alpha(t)dt = \begin{cases} \sqrt{1 + \left(\frac{\Delta D}{\Delta Y}\right)^2} \bullet \int_{Y_1}^{Y_2} \alpha(y)dy & (\Delta Y \neq 0) \\ (\Delta D) \bullet \alpha(Y_1) & (\Delta Y = 0) \end{cases} \tag{3}$$

This can be reduced further to two terms of a single variable,

$$\int_{Y_1}^{Y_2} \alpha(y)dy = \int_{-\infty}^{Y_2} \alpha(y)dy - \int_{-\infty}^{Y_1} \alpha(y)dy \tag{4}$$

which are pre-computed and stored in a one-dimensional lookup table. Thus, equation 1 can be evaluated in constant time with two table lookups. This value is used to blend the rendered surface color at each point with the color of the fog.

References
1 J. F. Blinn. Light reflection functions for simulation of clouds and dusty surfaces. In Computer Graphics (SIGGRAPH 82 Proceedings), July 1982, 21-29.
2 Marc Levoy. Volume rendering using the Fourier projection-slice theorem. In Proc. Graphics Interface '92, Vancouver, May 1992, Canadian Information Processing Society, 61-69.
3 Jackie Neider, Tom Davis, Mason Woo. OpenGL programming guide. 1993, Addison-Wesley, 216-221.

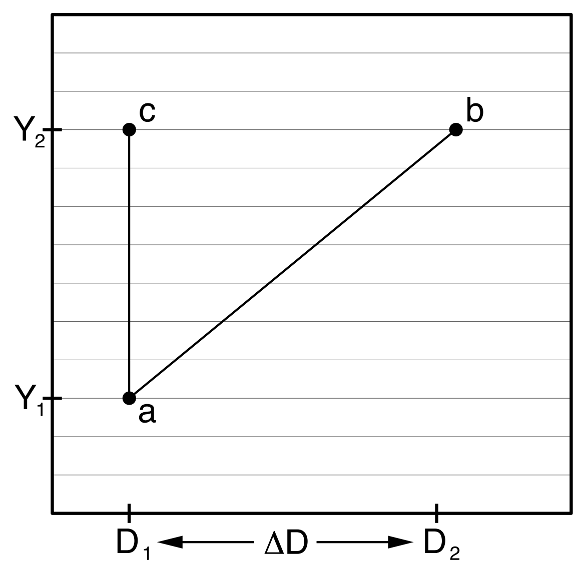

A Model for Anisotropic Reflections in Open GL

Wolfgang Heidrich
Computer Graphics Group
University of Erlangen
Am Weichselgarten 9
91058 Erlangen, Germany
heidrich@informatik.uni-erlangen.de

Many of the objects we know from day-to-day life have anisotropic surfaces. Often, this anisotropy is caused by a micro structure, in which long, thin features are aligned in one predominant direction. For example, brushed metal objects have been polished in such a way that there are a lot of parallel scratches in the surface. On CDs and records, the micro structure is given by the tracks carrying the music or data, and woven cloth consists of fibers oriented in certain directions.

The anisotropy in all these examples is created because the distribution of the surface normals along the scratches or fibers is different from the distribution across them. A model for this kind of surface has been described in by Banks.[1] The scratches or fibers are assumed to be lines or curves on the surface. With this assumption, existing lighting models for illumination of lines can be applied to rendering anisotropic surfaces. We are building on this work, and the work of Stalling, Zöckler, and HegeDetlev Stalling, Malte Zöckler, and Hans-Christian Hege,[2] for fast illumination of lines in OpenGL.

Illumination of Lines in 3D

Since our model is based on the assumption that the anisotropy of the material is based on microscopic scratches or fibers, we first discuss rendering of illuminated lines in 3D.

The fundamental difference between illumination of curves and illumination of surfaces is that every point on a curve has an infinite number of normal vectors. Thus, every vector that is perpendicular to the tangent vector of the curve is a potential candidate for use in the illumination calculation.

From this infinite number of potential normal vectors, the vector selected for lighting purposes is the projection of the light vector L onto the normal plane.[1]

OpenGL Implementation of Illuminated Lines

A line in 3D can be illuminated by using this normal vector in a standard Phong lighting calculation. As described in 2, both the diffuse component $I_{diffuse} = k_d \langle L,N \rangle \bullet I_i$, and the specular component $I_{specular} = k_s \langle V,R \rangle^n \bullet I_i$ can be expressed in terms of the cosines $\langle L,T \rangle$ and $\langle V,T \rangle$ only: the illumination at some point on the line is given as a bivariate function $I_o(\langle L,T \rangle, \langle V,T \rangle)$.

The basic idea of applying this illumination model in OpenGL is to code the tangent vector of the line segments into the texture coordinates. Then, the following texture matrix is used to calculate the two cosine values:

$$1/2 \begin{bmatrix} l_x & l_y & l_z & 1 \\ v_x & v_y & v_z & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 \end{bmatrix} \begin{bmatrix} t_x \\ t_y \\ t_x \\ 1 \end{bmatrix} = \begin{bmatrix} 1/2(\langle L,T \rangle + 1) \\ 1/2(\langle V,T \rangle + 1) \\ 0 \\ 1 \end{bmatrix}$$

When using these coordinates in combination with a precomputed texture map containing the values $I_o(\langle L,T \rangle, \langle V,T \rangle)$, the line is correctly illuminated according to the above model.

Anisotropic Surfaces and Self-Shadowing

The same basic idea can also be applied to rendering anisotropic surfaces. Here, however, we also have to take a certain self-shadowing term into account (i.e., regions pointing away from the light source should not be illuminated). In accordance with Banks,[2] we use the following model for dealing with self-shadowing:

$$I_o = I_{ambient} + clamp(\langle -N, L \rangle)(I_{diffuse} + I_{specular}),$$

where N is the actual geometric normal at the surface point, and the function "clamp" is zero for negative values, and the identity function otherwise.

This modification can be incorporated into the OpenGL implementation with the help of the standard OpenGL lighting mechanism. In the OpenGL rendering pipeline, the result of the texture lookup operation for each pixel can be multiplied by the result of the lighting calculation ("fragment color") for the same pixel. If we use a diffuse, white surface material, and a white, parallel light, the resulting lighting computations performed by OpenGL result exactly in the clamped cosine value required above.

Multipass Extensions

So far, we have only considered illumination through a single light source. Additional light sources can be handled by using multipass techniques. To this end, the object is rendered multiple times, once for each light source.

Another application of multipass techniques is the addition of isotropic reflection components. The model described above is purely anisotropic. Many surfaces, however, have both an isotropic and an anisotropic reflection component (for example, when the scratches or fibers on the surface have a certain, non-zero distance from each other. A single rendering pass with the standard OpenGL lighting model can add this isotropic component for all the light sources in the scene.

References

1 David C. Banks. Illumination in diverse codimensions. In Computer Graphics (Proceedings of SIGGRAPH 94), 327-334, July 1994.

2 Detlev Stalling, Malte Zöckler, and Hans-Christian Hege. Fast display of illuminated field lines. IEEE Transactions on Visualization and Computer Graphics 3(2), 118-128, 1997.

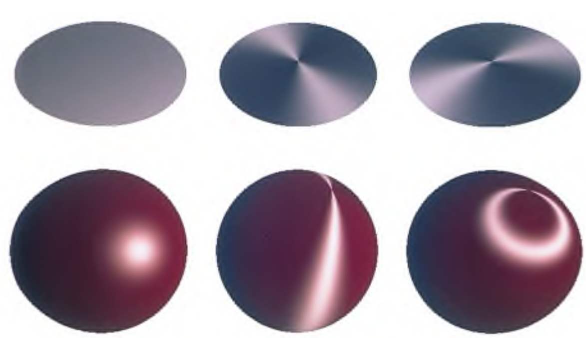

Quadratic Interpolation for Near-Phong Quality Shading

Larry Seiler
Mitsubishi Electric ITA
seiler@merl.com
www.merl.com/people/seiler

This sketch describes a simple way to set up coefficients to implement quadratic shading in rendering hardware. Quadratic shading interpolates a second-order shading function to provide quality comparable to per-pixel Phong shading,[1] at a fraction of the complexity and with no restrictions on the number of light sources or the type of lighting. Simple adaptive algorithms allow software and hardware to fall back to Gouraud shading.

Triangle Preparation
Fig. 1 shows a triangle with lighting computed at six positions: vertex colors L, M, and N, and edge midpoint colors Q, R, and S. Any lighting algorithm may be used to produce these colors. Edge midpoint colors need not be computed or transmitted to the rendering hardware if they are similar to the values produced by Gouraud shading. This may be determined by existing adaptive algorithms for subdividing Gouraud shaded triangles.[2]

Fig. 1 also defines the M and N vertex coordinates (x1, y1) and (x2, y2), as well as terms i = x2/x1, j = y1/y2, and k = 1–|ij|. The vertex with color L is translated to the origin and must be at a corner of the triangle's bounding box. As in the figure, assign the vertices such that i and j are no greater than 2, which is possible except for co-incident vertices. This reduces the numerical precision required for quadratic shading.

Quadratic Function Coefficients
We want to define a quadratic shading function $I(x, y) = ax^2 + bxy + cy^2 + dx + ey + f$ that evaluates to the specified color values at the six sample points. Equations (1) define five intermediate color values. Equations (2) define the six coefficients. T/2, U/2, and (T + U - 2V)/2 measure the worst-case difference between Gouraud shading and quadratic shading on the three edges, which occurs at edge midpoints. Rendering hardware may compare these to a threshold to select whether to fall back to Gouraud shading.

(1) $T = 3D\ L + M - 2Q \quad G = 3D\ 2Q - 2L - T$
$U = 3D\ L + N - 2R \quad H = 3D\ 2R - 2L - U$
$V = 3D\ L + S - R - Q$

(2) $a = 3D\ 2(T + Uj^2 - 2Vj)\ /\ (k^2\ x1^2)$
$b = 3D\ 4(2V - Ti - Uj - Vk)\ /\ (k^2\ x1\ y2)$
$c = 3D\ 2(U + Ti^2 - 2Vi)\ /\ (k^2\ y2^2)$
$d = 3D\ (G - Hj)\ /\ (k\ x1)$
$e = 3D\ (H - Gi)\ /\ (k\ y2)$
$f = 3D\ L$

The quadratic shading function may be easily computed using forward differencing. A scanline rendering algorithm requires two additions to compute each intensity. It can be defined such that stepping down the triangle edge requires four additions that also compute the first two intensities for the new scanline. Rendering with half-plane equations[3] requires three additions per intensity in the general case, since movement occurs along both axes.

Defined this way, quadratic shading does not require much more precision than Gouraud shading. Thirty-six-bit adders suffice for 8-bit intensity values within a 4K x 4K pixel address space, as compared to a 22-bit adder for Gouraud shading, including an overflow bit. Twenty-four-bit adders suffice for quadratic shading if long edges are subdivided so that x1 + y1 < 128 and x2 + y2 < 128.

Quadratic Shading Example
Figs. 2-4 graph the intensities computed for a 1 x 1 surface patch that is rendered with a light source at infinity above (0.7, 0.2) and with a specularity exponent of 10. The patch bends by 30 degrees in the X dimension and by 20 degrees in the Y dimension. Color bands indicate the level of intensity. Fig. 2 graphs per-pixel Phong shading on this patch. Fig. 3 graphs quadratic shading on two triangles using colors computed at nine points, three of which are shared between the two triangles. The result is quite similar to per-pixel Phong shading.

Fig. 4 graphs eight triangles that are Gouraud shaded using the same nine colors. This result produces artifacts such as Mach banding and a displaced, dimmer specular highlight. Animation exacerbates these problems. Eliminating the artifacts requires further triangle subdivision and many more Phong-lighting computations than quadratic shading requires.

This research was conducted at Digital Equipment Corp.

References
1 D. Kirk, D. Voorhies. The rendering architecture of the DN10000VS. Computer Graphics vol. 24, August 1990, 299-307 (see 303-304).
2 Y. Cho, U. Neumann, J. Woo. Improved specular highlights with adaptive shading. Proceedings of Computer Graphics International, June 1996, 38-46.
3 J. Pineda. A parallel algorithm for polygon rasterization. Computer Graphics vol. 22, August 1988, 17-20.

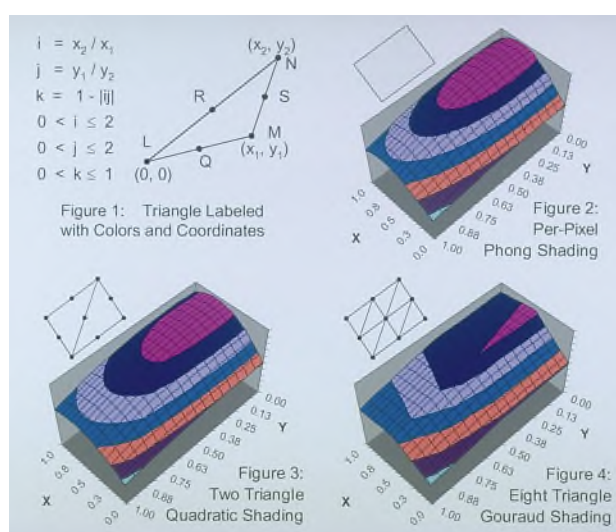

Figure 1: Triangle Labeled with Colors and Coordinates

Figure 2: Per-Pixel Phong Shading

Figure 3: Two Triangle Quadratic Shading

Figure 4: Eight Triangle Gouraud Shading

Editing 3D Objects Without 3D Geometry

Rui Yamada
Hashimoto Signal Processing Lab
Media Processing Lab
Sony Corporation
6-7-35, Kitashinagawa Shinagawa-ku
Tokyo, Japan
rui@av.crl.sony.co.jp

Mitsuharu Ohki
Hashimoto Signal Processing Lab
Media Processing Lab
Sony Corporation

Most conventional systems[1] that use image-based algorithms must make cumbersome preparations for reconstruction of 3D geometry of objects. Such preparations include acquisition of multiple images, calibration of camera(s), and specification of corresponding points in images.

We have developed a system that frees users from these preparation steps. Our system provides an easy-to-use 3D editing environment that requires only a single-source image to begin with. Users can edit 3D objects via their 2D projection on the image. A 3D object is interactively selected and exported to the paste buffer by specifying a couple of characteristic points of the shape. In the paste buffer, the object surface is flattened and laid out using a 2D coordinate that is naturally defined on the surface of the object. A user can perform edit and paint operations in this paste buffer. The contents of the buffer can be pasted either back to the original object or to another object.

The system also provides the following functions to assist users:

- **Points Adjustment**
  Once coarse positions of an object to be extracted are given, our system automatically refines these positions using the gradient of the image.

- **Diffuse Color Editing**
  Users can edit the diffuse (object) color while preserving the effect of illumination. The system performs a color histogram analysis to obtain shading and specular reflection in an image. Then, given a new diffuse color, the system adjusts its intensity to match the shading in the original image and add the specular component to restore shiny reflection. The same method is used in our illumination conscious texture mapping.

- **Spreading out Adjacent Surfaces to One Paste Buffer**
  Adjacent surfaces in the image preserve the adjacency in the paste buffer as well. This is useful when a user wants to draw across surfaces.

- **Image Restoration**
  The system can remove an unwanted foreground object (i.e., 3D eraser). This is done by applying an image restoration algorithm[2] in the paste buffer.

- **Alpha Blending**
  Optional key (alpha value) of the destination image can be used during the paste operation. Users can write-protect regions of the destination from overwritting by the paste buffer contents, achieving additional occlusion effects on the result.

Fig. 1 shows a sample of an editing session using our system. In the result image, the covers of the yellow and purple notes were swapped and the cover of the yellow note was pasted on the stuffed cat using the alpha blending. The textures of the green ball, the bottle, and the white wall were changed. The blue paper and the CRT body were painted. The original image was pasted on the CRT. The hole on the brown wall was repaired.

Our system is general enough to handle various 3D objects provided that the user can give the shape hint. This is a reasonable assumption, especially for interior scenes in which most objects are industrially designed from simple shapes.

References
1 Seitz, S. M. and Kutulakos, K. N. Plenoptic image editing. In Procceedings of ICCV, 1998.
2 Hirani, A.N. and Totsuka T. Combining frequency and spatial domain information for fast interactive image noise removal. In Proceedings of SIGGRAPH 96, 1996, 269-276.

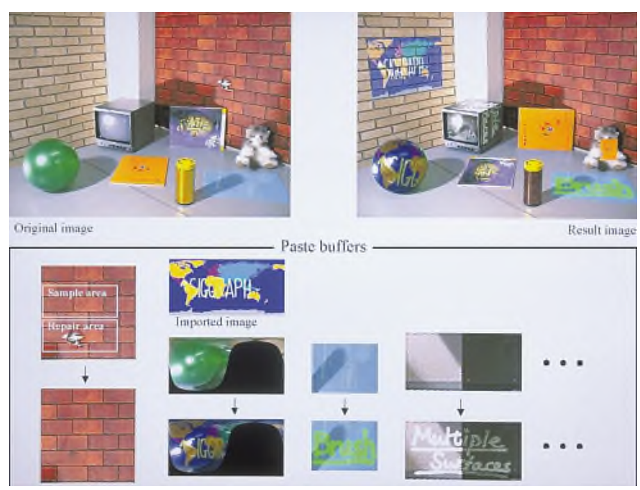

Figure 1 A sample of editing session

Procedural Field Grasses

Lee A. Butler
Attn: AMSRL-SL-BV
U.S. Army Research Laboratory
Aberdeen Proving Ground, Maryland 21005-5068 USA
butler@arl.mil

David S. Ebert
University of Maryland, Baltimore County

One of the greatest challenges remaining in computer graphics is modeling natural environments, as such scenes contain a staggering amount of detail. Practical limitations on both time and storage make it impossible to describe scene elements such as grass in geometric detail. Texture maps can help, but they present their own burden in selection and management. Ideally, it should be equally easy to specify an object as "grass" or "mirror" or "plastic." Thereafter, the object should exhibit the appropriate visual behavior with minimal human intervention.

Observations

Unlike most plants, which grow at their tops, grasses grow at the center of their base. Growth pushes older parts outward and upward. At maturity, they produce an upright central stalk. In images containing grasses, the green grass blades dominate the near field. The effect of individual grass blades can be seen from a moderate distance, as sunlight reflects off their relatively flat surfaces. As distance increases, the optical density created by sparse central stalks from plants and their taller blades becomes substantial.

Previous Work

Other researchers have developed methods for modeling grasses. Smith[1] presented a method called "graftals" for producing individual plants at varying resolutions using a formal grammar. Populating a terrain with such plants is tedious. Kajiya and Kay[2] constructed voxelizations of representative samples of 3D textures. These "Texels" are tiled over a surface. Some effort must be expended to blend the seams between voxels. Texels also tend to produce grid artifacts when rendered from a distance. Proper resolution for the voxels can be difficult to determine. Goldman[3] demonstrates optimizations for rendering fur to be viewed from a relatively large distance. The visual behavior of fur and grass are similar under these conditions

Method

The desire for vast tracts of grass, easily specified and self-animating, strongly suggests a procedural approach to modeling. In our model, the user specifies a volume to be occupied by the grass, a terrain grid cell size, and the mean and variance for the number of grass plants per cell. Grass plants in the volume are generated from a prototype as they are needed at rendering time. The prototype consists of a list of blades emanating from a single root. A blade is defined by a width and a number of linear leaf segments, each of which in turn consists of a vector, length, and surface normal. The central stalk is represented as a blade with minimal curvature and a different color. To make the process of creating plants from the prototype efficient, the list of blades is ordered. Young or most upright blades come first in the list, followed by older or more curved blades, with the central stalk last. A single random number allows us to select the number of blades and a scaling factor for their length.

The lower surface of the volume acts as the terrain on which grass is growing. At rendering time, a ray cone is marched through the cells defined by the terrain grid. Cells adjacent to those that the ray passes through are also evaluated. This allows plants rooted in one cell to extend into the adjacent cell, which greatly reduces grid artifacts. A random-number generator is called with the cell coordinates as the seed to obtain the number of grass plants and their location in the cell. The prototype plant is tailored by a number of noise functions to vary the height, number of blades, and rotation. Turbulence noise is applied to generate the effects of wind. It is possible that a large number of cells will have to be interrogated before finally finding a plant intersection, especially at low grazing angles, making rendering times for such areas overly long. Since there is little or no detail content in such areas, the expenditure is not justified. It is possible to recognize this condition and substitute a further approximation. We invoke our approximation when the ray diameter becomes substantially larger than the maximum grass blade width and approaches the size of a grid cell. At this point, a noise function produces a surface normal. A surface color is chosen by using the altitude of the ray-cell intersection and depression angle of the ray to interpolate the green of blades and the brown of stalks. Experimentation has shown that linearly interpolating RGB colors based upon a linear scale of altitude with depression angle produces acceptable results.

Conclusions and Future Work

We have shown that it is possible for complex grass structures and behavior to be easily generated with minimal storage requirements. Convincing foreground detail and motion can be achieved, and our adaptive algorithm reduces computational costs. One possible area of extension is to incorporate multiple prototype plants.

References

1 Smith, Alvy, Ray. Plants, fractals, and formal languages. Computer Graphics. 18,3 1-10.

2 Kajiya, James T., and Kay, Timothy L. Rendering fur with three dimensional textures. Computer Graphics, 23,3 271-280.

3 Goldman, Dan B. Fake fur rendering. Computer Graphics, 31,3 127-133.

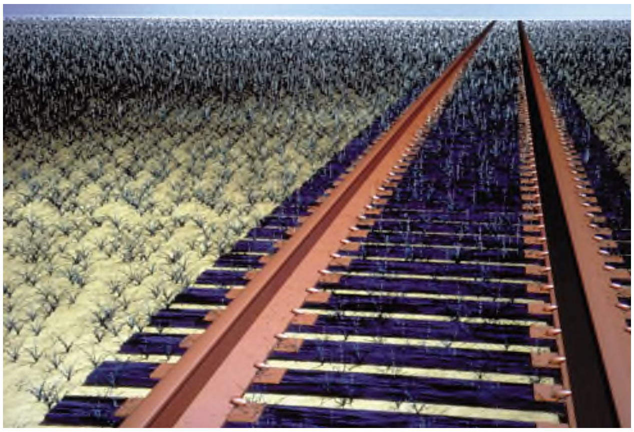

Modeling Fiber Stream of Internal Wood

Naoki Kawai
Communication & Information Laboratory
Dai Nippon Printing Co.,Ltd.
kawai@lab.cio.dnp.co.jp

Generating the texture of wood grain is a lot of fun and sometimes useful as well.

One of the most obvious features of lumber surfaces is the annual rings. They are relatively easy to picture using computer graphics, and they can be modeled as coaxial cylinders in a 3D world coordinate system. Then the image can be generated by projecting a cross section onto the screen.[1] Deforming the cylinders using a fractal technique is known to be effective in making the image more realistic (Fig.1). However, the generated pictures are not sufficiently convincing in terms of photo-realistic expression.

Annual rings are not the only element that constitutes the appearance of wooden surfaces. Another feature is the lack of uniformity in reflection called "figure." This is continuously changing depending on the continuous changes in the combination of the direction of the light, the angle of the surface, and the viewing position. An example in which this feature is best utilized is the back of a violin (Fig.2) .

It is well-known in forestry that the figure is caused by the cross grain,[2] a certain pattern formed by the wood's internal fibers, which are usually not straight.

Modeling Cross Grain

In this section, some types of cross grain are explained and a method of modeling them is introduced.

In general, the fiber elements in trees are not arranged strictly along the longitudinal growth axis. The irregularity in the structure is called cross grain. Fig.3 shows three typical patterns of grain: straight grain, spiral grain, and wavy grain.

Grain patterns are defined in a cylindrical coordinate system (τ, Θ, y), in which y represents the central axis of the wood in the growth direction, and τ and Θ represent polar coordinates on the plane perpendicular to the y axis.

Straight grain's structure is the easiest one to define mathematically because the structure is expressed independently of y. Spiral grain and wavy grain can be represented as partial transformation in Θ component from straight grain in the following formulae :

Spiral grain: $\Delta\Theta = \delta y / \tau$
Wavy grain: $\Delta\Theta = A \sin(2\pi y / \lambda)$

where δ is the gradient of spiral, λ is the wave length, and A is the amplitude of wave. Any of these grains can be combined arbitrarily, and the parameters can be defined as functions of y, i.e., $\delta(y)$, $\lambda(y)$ and $A(y)$ with additional randomness using fractal techniques.

Reflection Depending on Cross Grain

In this section, the way the cross grain influences reflection on the lumber surface is explained, and a simple technique to verify it is shown. Since fiber cells are formed roughly as tubes, a lumber surface is covered with a large number of microscopic cylindrical grooves. These grooves are known to make the reflection anisotropic.[3] Furthermore, the bottoms of these grooves are typically inclined along the running directions due to the cross grain. The inclinations vary by location of the lumber surface, which causes both a lack of uniformity and a changeability in reflection.

A simple, though not precise, way to realize this phenomenon is to manipulate the normal vectors by bump-mapping the inclinations to the lumber surface when calculating reflections on it.

Results and Future Works

Fig.4 shows the lumber surface with changeable reflection generated by the above-mentioned techniques. These four pictures show an identical lumber surface with differences in lighting direction. Since cross grain is defined in 3D space continuously, it is also possible to render a curved surface with changeable reflection.

Our current interest is to model other features of wood such as rays (radial tissues). Furthermore, a precise model of anisotropic distribution of reflection may improve the quality of generated images.

References
1 Peachey, D.R. Solid texturing of complex surfaces. Proceedings of SIGGRAPH 85, August 1985.
2 Ertl,J. Die anatomie des holzes. Das Grosse Buch vom Holz, 1977.
3 Poulin, P. and Fournier, A. A model for anisotropic reflection. Proceedings of SIGGRAPH 90, August 1990.

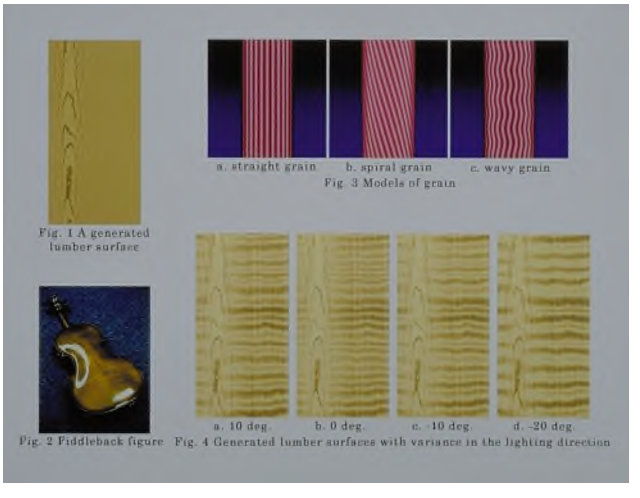

Fig. 3 Models of grain
Fig. 1 A generated lumber surface
Fig. 2 Fiddleback figure
Fig. 4 Generated lumber surfaces with variance in the lighting direction

Session: Natural Causes

Botanical Tree Structure Modeling Based on Real Image Set

Tatsumi Sakaguchi
ATR Media Integration & Communications Research Laboratories
+81.774.95.1433
+81.774,95.1408 fax
tatsu@mic.atr.co.jp
www.mic.atr.co.jp/~tatsu/

Some studies of botanical tree modeling have been based on the growth model, which can construct a very natural tree structure.[1] However, this model makes it difficult to predict the final form of the tree given parameters; that is, if an objective form of a tree is given and it is to be reconstructed into a three-dimensional model, we have to change the parameters to reflect the structure by a trial-and-error technique. Thus, we propose a new top-down approach in which a tree's form is defined by volum data that is made from a captured real image set, and the branch structure is realized by simple branching rules.

In this technique, some still color images of the target tree are first acquired with a video camera from different viewpoints, and unnecessary scene background is removed by manual operation. Image pre-processing involves four successive stages: normalizing trunk position and the tree size between each image, analyzing the color texture of branches and leaves, generating the color textures for the target tree, and color-to-gray-scale image conversion. A volume data construction process using the inverse reprojection method is performed for the gray-scale image. The volume data are defined as an object that fully contains the target tree and consists of voxels, each with a density value that corresponds to the pixel value in the original image.

Simple branching rules then create the tree-branch structure within the volume data. The rules include a simple ternary branch model and some restrictions for expressing a general broadleaf tree. More specifically, a branch is formed by connecting branch segments that have cylindrical shapes, and a branch can split into a maximum of three directions. Special rules are not needed for each kind of tree because the form differences of various species can almost be represented from the volume data. The branching rules are controlled by the density distribution of voxels in the volume data and restrictions on the tree structure in creating a natural form. A branch basically extends in the direction of the largest density value in the volume data. For example, thick branches, a trunk, or an area thickly covered with leaves gives a large density value. It never extends to outside of the volume data. In addition, we apply some restrictions on the branching rules: avoid generating a hook-shaped branch, avoid generating a branch extended towards the ground, make branches extend from the trunk position as much as possible, and give higher priority to a straight branch.

Finally, the predefined leaf models are attached to the branch structure model, and the color textures obtained from the target image are mapped onto each polygonal model. We can create a botanical tree model that has a form much like the target tree. Since the tree model made by our method has a branch structure, its dynamics can be applied to a natural environment simulation system.

Reference

1 R. Mech and P. Prusinkiewicz. Visual models of plants interacting with their environment. SIGGRAPH96 Computer Graphics Proc. Annual Conference Series, 397-410, 1996.

A Particle System for the Direct Synthesis of Landscapes and Textures

Argiris A. Kranidiotis
University Of Athens
Alex. Panagouli 32
Keratsini 187 56 Greece
akra@di.uoa.gr

We present a new fast method for generation of landscape images and textures. A particle system following simple rules generates images that give the illusion of a 3D model, although images are generated directly in 2D space. By varying the parameters of the system, the user can generate a plethora of different images ranging from snowy mountains to marble textures.

Traditionally, fractals were used successfully for realistic generation of mountains and landscapes.[2,8,5,3] In contrast, our method does not use 3D polygons. In fact, the image is generated directly in two dimensions. Images produced by our algorithm look remarkably natural. In addition, the whole process is computationally efficient, and images are generated in seconds. Still images of this type are required in a wide range of applications: automatic background generation for traditional animation, texture synthesis, pen-and-ink illustrations, and even as a source of inspiration or experimentation for artists.

The basis of our algorithm is a particle system.[6,7,4] In our system, particles move according to simplified laws of physics and some additional heuristics. Note that our purpose is rapid generation of a convincing image, rather than computing an expensive accurate physical simulation.

Description of the System

Mountains are generated when many particles accumulate together. Particles are generated by many sources that coexist at the top of the scene. The simulation is terminated only when mountains reach a user-defined height. We consider particles that move inside a liquid according to Stokes approximation: $T = 6\pi\eta \cdot r \cdot u_\infty$.[1] Other forces that influence their motion are gravity, resistance of fluid, and collisions between static particles

The main concept behind the behavior of the particles during a collision is that particles should become immobilized when they reach a flat surface, but they should roll downward when they hit an inclined surface. Due to the discrete nature of the system, the surfaces formed by the immobilized particles are not continuous. Using simplified rules, our system classifies the slope of the surface at a point as: flat, positive, negative, and undefined.

When a particle hits the surface at some point, then its behavior is defined as follows:

- When the surface hits a flat surface, the particle stops its motion forever.
- When a particle hits the surface at a positive or negative slope, it moves to the free pixel pleft or pright respectively, and finally it continues its motion with a new random-direction vector $\vec{d}_r$.
- If the surface slope is undefined, then it can move with a 50-percent probability to either pleft or pright, and then it continues its motion with a random new-direction vector $\vec{d}_r$.

Our system consists of two distinct types of particles, each with different colors and collision statistics. We call one type black and the other white. Black particles represent the mountains, while white particles represent the snow. In addition, the mixture of these two creates interesting shadowing and highlight effects in the image. The behavior of the particles during a collision is crucial for creation of realistic images. Snow tends to accumulate in the cavities and slopes of a mountain. To achieve the same effect, white particles move significantly faster in the x axis when they bounce off.

All particles in the system are generated by discrete moving sources that exist at the top of the scene. Each particle source has its own statistical characteristics and parameters that vary inside some user-defined range. Moreover, sources can overlap each other.

Our algorithm has also proved successful in the field of texture synthesis. If we increase the speed range of the sources by a factor of about 50, particles can not form mountains. Instead, they create images that closely resemble marble or sedimentary rock.

Using super-sampling techniques, we can create grayscale images. Even more realistic images can be generated by direct mapping of the gray-scale levels to another colormap (Fig. 1).

Conclusions

The potential of our method has not been fully exploited. Further experimentation with the system parameters will certainly produce new unpredictable results, new rich textures, etc. Other types of colored particles can be easily incorporated in our system to simulate the effects of verdure, different types of ground, lakes, etc. Also, we believe that the same model can easily extend in three dimensions, in order to generate a polygonal mesh. As an added benefit, the algorithm will also produce a texture map of the snow, as well as the 3D model at the same time.

References

1 Alexopoulos, K.D. General Physics, Volume I, Athens 1960, 247.
2 Fournier, A., D. Fussel, and L. Carpenter. Computer rendering of stochastic models. CACM, 25(6), June 1982, 371-384.
3 Miller, S.P., Gavin. The definition and rendering of terrain maps. SIGGRAPH 86, 39-48.
4 Xin Li and Moshell M. Modeling soil: Realtime dynamic models for soil slippage and manipulation. SIGGRAPH 93, 361-368.
5 Peitgen, H.-O. and Saupe eds. The Science of Fractal Images, Springer-Verlag, New York, 1988.
6 Reeves, W.T. and R. Blau. Particle systems – a technique for modeling a class of fuzzy objects. SIGGRAPH 83, 359-376.
7 Reeves, W.T. and R. Blau. Approximate and probablistic algorithms for shading and rendering particle systems. SIGGRAPH 85, 313-322.
8 Voss, R. Random fractal forgeriesl, in Ernshaw, R.A., ed., Fundamental Algorithms for Computer Graphics, Springer,-Verlag, Berlin, 1985; NATO ASI series F, volume 17, 805-835 .

Figure 1 Image rendered directly by our particle system. Although the simulation was strictly 2D, the image gives the illusion of a 3D landscape.

Session: Natural Causes

Plane-Shape Perception Using Point-Contact Type Force Feedback Device

Juli Yamashita, Yukio Fukui, Osamu Morikawa
National Institute of Bioscience and Human Technology
Tsukuba 305-8566 Japan
{ juli | fukui | ycai | morikawa}@nibh.go.jp

Cai Yi
Japan Science and Technology Corporation

Robert W. Lindeman
Department of EE & CS
The George Washington University
Washington, DC 20052 USA
gogo@seas.gwu.edu

Force (haptic) feedback devices (FFDs) form an important field of research in virtual reality. They are especially promising in applications that require both three-dimensional input and output, including CAD systems and surgical simulators. Haptic rendering requires a very high control rate (~kHz), while CAD modeling and physics-based simulation processes are often 100 times slower; to bridge the gap, haptic rendering must be separated from other processes. We propose feature-based haptic rendering and its communication protocol[1] to exchange shape information as features (shape fragments, haptic characteristics, etc.) between processes (Fig. 1).

Modeling/simulation selects and transmits features, which are haptically rendered using an FFD. Here, the latency gap between processes requires feature interpolation to smooth haptic rendering. Features can be interpolated in both space and time domains; spatial interpolation is important for smoothing polygonal surfaces, and temporal interpolation is necessary in shape deformation and touching moving objects. To develop less expensive but effective interpolation algorithms, we must know much more about human haptic characteristics. A modeler process, for example, can prefetch features before time-consuming work and human sensory thresholds define their resolution. We conducted experiments on static plane-shape recognition thresholds using the PHANToM2 (SensAble Technologies) as an FFD.

Experiment 1

Subjects (11 right-handed Japanese adults) touched a stimulus shape (Fig. 2) using a stylus probe and altered its height, h, with an up or down arrow key. Reaction magnitude f was calculated as $f = (G(x) - y) \cdot s$, where x, y are the coordinates of the stylus tip measured in mm, and s = 0.5 is stiffness. The reaction force direction was fixed upward, (0, 1, 0), the same as the normal of a flat plane, because varied reaction force directions cause haptic illusion.[3] In the ascending experiments, the initial stimulus shape was a "flat" plane with a small height that could be increased by pressing the up arrow key until the subject felt it to no longer be flat. In the descending experiments, a mountain shape was initially presented and the subject decreased h with the down arrow key until it felt flat. In both cases, the subject hit the return key to signal the final h.

Three trials were conducted for each ascending/descending for 5 width w (5, 10, 20, 30, and 40 mm). The h changed by 0.5 mm (w = 20, 30, 40 mm) and 0.1 mm (w = 5, 10 mm) each time arrow keys were pressed. To avoid the subject's remembering the number of key presses, initial h value was altered in each trial. The order of trials was also randomized.[2] To determine the effect of stiffness, a preliminary experiment was conducted with five subjects, in which stiffness s = 0.25 and width w = 10, 20, and 30 mm.

Results and Discussion

The median of the three h obtained under each condition was used to compute the average of ascending and descending experiments for each w. Fig. 3 shows fh, the force for height h ($fh = h \cdot s$), and Fig. 4 shows force gradient fh/w. Thin black lines are experiment 1 results and thick red lines experiment 2 results.

Experiment 1: The resulting fh for w = 5 and 10 mm do not differ statistically (one-way ANOVA: $F(1, 20, 0.05) = 4.35 > F0 = 0.06$), implying that human force sensitivity saturates at about 0.2 - 0.3 N. With widths larger than 10 mm, fh/w is almost constant; a one-way ANOVA gives $F(3, 40, 0.05) = 2.84 > F0 = 0.75$, indicating that, when the direction of the reaction force is fixed, a shape will be felt as flat if the force gradient is 0.02 - 0.03 N/mm or less ($w \in [10, 40]$).

Experiment 2: Despite the small number of subjects, results show a tendency similar to experiment[1], except that force sensitivity saturation appears to start at a w of 10 - 20 mm. Further research is needed on the effect of the stiffness parameter.

Conclusion and Future Work

Thresholds of human haptic shape perception using the PHANToM device are: (1) Force gradient: 0.02 N/mm and (2) Minimum force: 0.2 N. Although the curvature effect requires further study, a curved surface should also be perceived as being smooth if this force gradient is satisfied. Future work includes investigating the threshold with smaller width, the effects of curvature and stiffness, and temporal perception characteristics.

References

1. Yamashita, J., Y. Cai, and Y. Fukui: Feature-based haptic rendering – architecture, protocol, and application. Visual Proc. SIGGRAPH 97, 1997. (Revised version available at: www.aist.go.jp/NIBH/~b0621/Project/FBR97.ps.gz)
2. Salisbury, K. et al. Haptic rendering: programming touch interaction with virtual objects. Proc. ACM 1995 Symp. on Interactive 3D Graphics, 1995.
3. Minsky, M. et al. Feeling and seeing: Issues in force display. Computer Graphics, Vol. 24, No. 2, 1990.

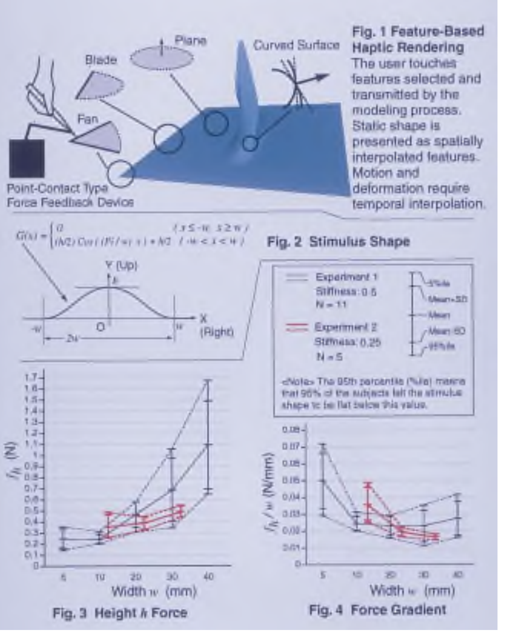

Fig. 1 Feature-Based Haptic Rendering The user touches features selected and transmitted by the modeling process. Static shape is presented as spatially interpolated features. Motion and deformation require temporal interpolation.

Fig. 2 Stimulus Shape

Fig. 3 Height h Force

Fig. 4 Force Gradient

A Psychophysical Study on Human Perception of Surface

Mikio Shinya
NTT Human Interface Laboratories
shinya@nttcvg.hil.ntt.co.jp

Shin'ya Nishida
NTT Basic Research Laboratories
nishida@apollo3.brl.ntt.co.jp

Human perception is an important element in image synthesis, and its better understanding can lead further advances in computer graphics. Computer graphics technologies, on the other hand, can provide powerful experimental tools to psychologists. This sketch shows a successful example of collaboration among psychologists and computer scientists in an attempt to analyze human perception of surface reflectance properties.

Surface reflectance properties can be represented in the form of bidirectional reflectance density functions (BRDFs). Although a BRDF well describes the physical properties of surface reflection, the question addressed here is how precisely human beings can identify BRDFs. We examined the possibility that the human visual system extracts some "filtered" information from reflection properties, just as we perceive color through trichromatic sensation rather than the spectral distribution itself.

Previous Work

The intensity of reflected light depends not only on BRDFs, but also on many other factors such as direct/indirect illumination and the geometric structures of the surface. Thus, it seems a difficult task to extract surface properties from an image sequence. However, previous studies have revealed some of the smart abilities of our visual system. It is known that reflectance perception is fairly stable under changes in illumination, referred to as lightness consistency. Also, Gilchrist and Adelson demonstrated that reflectance perception takes account of global geometric structures in some simple cases. This sketch investigates the influence of geometric factors on reflectance perception in more general situations.

Experiment

We examined whether subjects could identify the same reflectance property on different geometric surfaces. As in the Image Plate, two rendered surface (target and match) were displayed side by side on a monitor, and both surfaces were rotated along the horizontal axis to provide 3D cues. The task was to interactively choose the reflectance parameters of the match surface so that the reflectance properties of the two surfaces appeared as similar as possible. As the reflection model, we adopted the Phong model.

Results and Discussion

Five subjects were tested with a variety of surface shape combinations. If the human visual system is smart enough, we would expect the reflectance parameters of the target to differ only slightly from those of the match surface providing the best appearance match (constant errors). In most combinations, however, we observed large constant errors. This demonstrates that our visual system cannot properly take into account the geometric structures in establishing reflectance perception.

One of the subjects reported that he might have conducted the matching task based on image features, such as the size of highlights, average luminance, and so on, rather than reflectance properties in terms of BRDFs. To check this idea, we calculated the intensity histograms of the target and match images, and analyzed their differences. It turned out that most matched parameters provide small histogram differences between the target and match, close to the minimum distance. On the other hand, the histogram difference when the match image is rendered with the same parameters as the target is much larger in most cases. This supports the idea that we simply chose the parameters that provide similar images, neglecting the difference in surface geometries.

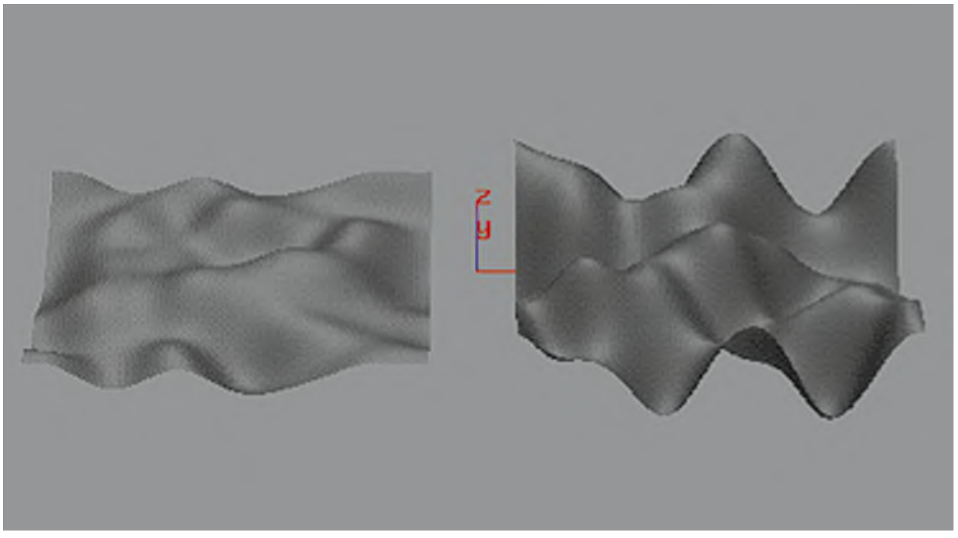

The Analysis and Visualization of Metamorphopsia through 3D Scene Regeneration

Peter Presti, Stephen Sinclair, Amit Gupta
Interactive Media Technology Center
Georgia Institute of Technology
Altanta, Georgia 30332-0130 USA
peter.presti@oip.gatech.edu

Metamorphopsia is a vision defect caused by a distorted retina in one or both eyes. Affected patients perceive this condition as a dynamic distortion in the geometry of their environment. Metamorphopsia transforms the patient's world into a perpetual Escher-like environment. Mild horizontal distortions produce binocular depth perception errors as evident in the attached figure. More severe horizontal distortions as well as vertical distortions cause ghosting or frank double vision. The patient finds navigating stairs or non-uniform topologies a challenge. Sudden onset of the distortion causes the greatest difficulties, whereas if the distortion develops more slowly, there may be some adaptation to the sensory input such that the patient may not even realize that a visual defect exists.

A tool to document metamorphopsia has been developed as part of the Multi-Axis Vision Evaluation System (MAVES). With distortion data derived from testing a patient with MAVES, the 3D environment can be reconstructed from the patient's point of view. This process involves projection of the 3D dataset onto each binocular 2D retinal plane. The planar data are perturbed by the documented distortions, and two 3D datasets are reconstructed. Visualization of the reconstructed datasets allows a third party, such as the ophthalmologist or family member, to observe the perturbations the patient sees.

The MAVES testing apparatus consists of a 21-inch video monitor and PC, a chin-rest to fix the patient's head, an interaction control pad, and fixation monitoring via an infrared pupil tracker. The metamorphopsia analysis program presents 5x5 regular spline grid with random perturbations overlaid on a cross, which the patient fixates at its center. The display grid dims if the pupil tracker detects a fixation deviation. The patient is instructed to remove the distortions from the grid until all of the lines appear straight and uniform. This is accomplished by manipulating over the grid using buttons on the control pad and adjusting the deviation of the lines using a joystick. The resulting grid objectively documents the macular distortions.

The distortion data from the resultant grid may be applied to a flat, 2D image. The distortion maps the image from the world image space to the patient's image space. An image with the distortion applied will appear geometrically correct when presented to the patient. The inverse distortion maps the image from the patient's image space to the world image space. An image with the inverse distortion applied is representative of the patient's monocular vision aberration and has been extremely useful in demonstrating to the ophthalmologist or family the monocular distortion perceived by the patient. However, the 2D image fails to adequately express the real-world 3D distortion the patient perceives with binocular vision.

A CG 3D environment may be reconstructed to demonstrate what is perceived by the patient under binocular conditions. This process involves projecting the 3D point coordinates of the dataset onto each binocular 2D retinal plane through each eye's nodal point. The 2D planar points are perturbed by the documented distortion measured for each eye. Two vectors defined by the perturbed 2D points in each retinal plane and corresponding nodal point are projected back into the world space occupied by the original geometry to create two new 3D geometry datasets. The position where the two vectors are closest is analyzed. If the distance between the vectors (d) is smaller than a convergence error limit (k), then the center point of the line drawn between the two vectors defines the new 3D point for both geometry data sets. If $d>k$, then two new 3D points are defined by the point of closest approach for each vector.

The two 3D geometry sets are rendered to produce two images, one for each eye. A composite of both images is generated with 50-percent transparency to produce a single image representative of the patient's environment. As the patient moves, new geometry sets are calculated and rendered.

The warped 3D geometry appears to realistically represent the visual disturbances that these patients describe and for the first time has given clinicians an insight into this poorly understood and poorly appreciated aberration of vision.

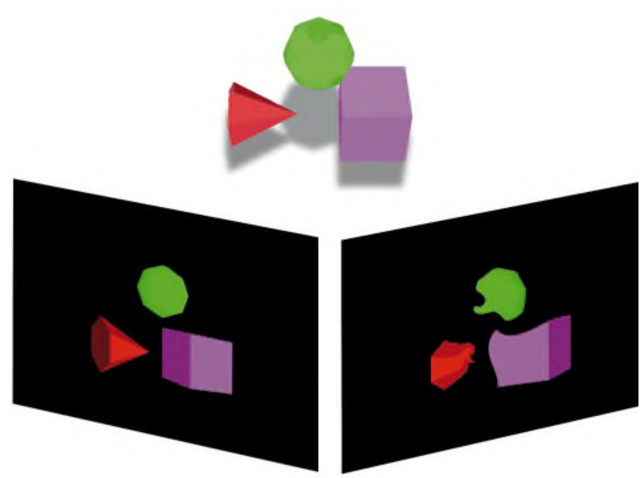

The Use of Distortion for Special Venue Films

Toshi Kato, Keith Goldfarb, Bob Powell
Rhythm & Hues Studios
5404 Jandy Place
Los Angeles, California 90066 USA
bpowell@rhythm.com

One area of computer graphics that has blossomed in recent years is special venue ride films for theme park attractions. The venues for these films usually have very specialized requirements for screen size, orientation, and projection methods. These specialized venues pose unique technical problems for an animation. How does one create an image that, when projected onto a dome or the inside of a torus, will appear undistorted to the viewer?

This sketch discusses three real-life productions from Rhythm & Hues Studios where distortion of images for these venues had to be implemented. Each of these projects presented unusual technical challenges, and each was solved in a different way. This sketch outlines the initial problems each venue posed and how we solved those problems for the final production. It then summarizes the lessons learned in each project, so they can to be applied to future projects.

The Kia Project

The screen for the Kia job was irregular. It was roughly shaped like the interior of a torus (i.e., it curved toward the viewer left-right and also up-down, similar to a piece of a spherical screen, but with a more extreme curve.

However, the overall field of view was similar to that of an IMAX projection, approximately 90 degrees horizontally and only half that vertically. A "mock-up" screen and projector were available to us nearby, so we could experiment in a theater that was very similar to the actual theater in size and shape. The use of such mock-ups has proven invaluable.

We chose to render from a single camera and use a post-rendering distortion. Since no algorithmic information about the lens nor the screen was available, we couldn't algorithmically determine the distortion. As a result, we needed to gather our own data empirically.

First, a standard "12-field" grid chart was projected in the theater and used to measure the angle and elevation of the grid points as viewed from a "best-seat" position. The measured angles and elevations were not evenly spaced, due to the lens and screen distortions. This spacing of grid points gave us a mapping of our film space into viewer space that we could invert. Then, using Mathematica, we found a "best-fit" polynomial for the distortion mapping and the desired corresponding inverse function.

The last step was to sample this inverse function at evenly spaced intervals. We made a piece of film with dots placed at points determined by this function, and when we projected it in the theater, these dots appeared all aligned and evenly spaced. This verified our inverse function. Then we used these points as input to a morph program, so the rendered images, after morphing, looked correct when viewed in the theater.

The Hilton Hotels "Star Trek the Experience" Project

Hilton's "Star Trek the Experience" project had a dome screen. Here, we knew both the size of the final screen, the projector location, the best seat, and the type of lens the projector would use. A post-distortion algorithm was chosen because we wanted the ability to pan and zoom the images as a 2D process.

One of the most interesting aspects of the distortion for Trek is that three camera images were used to generate the final image. No one camera would have had a wide enough field of view. Thus the center camera, or sweet spot, was used to generate a high-resolution 2K image, and the other two cameras used a lower resolution, which sped up the rendering time. Then the distortion algorithm seamlessly blended the images from these three cameras together.

A mock-up dome allowed us to project and test the distortion algorithm in-house.

The Caesar's Palace "Race For Atlantis" Project

The Caesar's Palace project was unique. Although it is also projected on a dome, this project required a stereo projection that would be viewed by the user at any angle, since the dome encompasses the field of view. Due to the curved nature of the dome surface and the stereo requirement, neither a post-distortion technique or the standard two-camera stereo technique would work .

Once again using information about screen dimension, lens distortion, and the best seat position, we used a backwards ray-tracing technique to map the geometry into the dome space. First all scene information was converted to dome coordinates. Then for every vertex in the scene, the vertex was projected onto the dome screen based on the point of view of an ideal user sitting in the best-seat position. Then using the lens distortion curve and the physical projector setting, we were able to calculate the vertex's projection onto the film plane. This calculation replaced the standard "world-to-screen" coordinate mapping in the renderer. Then the Rhythm & Hues renderer was rewritten to handle the special cases of tessellating all the geometry in the frame based on its mapping into this distorted screen space and to do specialized clipping of the tessellated geometry in this space. One of the key issues for this distortion was taking into account the user's head rotation. The distortion mapping changes if the user is looking up or to the side, rather than straight ahead, so the amount of distortion in these cases had to be smoothly interpolated across the surface of the dome.

Hairstyle Simulation System

Kazunori Miyata
Department of Imaging Art, Faculty of Arts
Tokyo Institute of Polytechnics
2-9-5 Honcho, Nakano-ku
Tokyo, 164-8678 Japan
miyata@img.t-kougei.ac.jp

This hairstyle simulation system automatically fits a stored hairstyle image to an input human face image.

To match various hairstyle images with a person's face by using photo-retouching software, it is necessary to modify the images by adjusting their size, location, and angle manually, and then paint in the gap between the hairstyle image and the face image. This retouching process may take about half an hour for each hairstyle, and is very tedious for the user.

Alternatively, it is possible to generate hairstyle images with computer graphics techniques, but this approach is impractical in view of the large data volume, the long processing time, and the low level of realism of the generated images.

Our system has several advantages over similar existing applications. For example, other systems require preprocessing, such as image resizing and background image removal for each captured image, whereas ours does not. In addition, other systems do not take account of gaps between the face and hairstyle, whereas ours fills the gaps by deforming the hairstyle image.

In our method, the selected hairstyle image is composed with a captured face image in a few seconds through simple operations, and the result does not appear artificial. In addition, the user can change the length of the composed hair easily and freely.

The images, which are all two-dimensional and in color, are input from a digital camera. The face boundary and the positions of facial features such as the eyes and jaw are specified manually with a pointing device, such as a mouse, or extracted automatically with an image recognition process.

The target hairstyle image is adjusted in accordance with these data as follows. First, the widths of the faces in the captured face image and the hairstyle image are calculated, and the hairstyle image is then scaled accordingly. The angle of facial tilt is calculated for both the face and hairstyle images, and the hairstyle image is then rotated by the difference of these angles. Finally, the hairstyle image is shifted by the location difference of the reference point, which is the center of gravity of the triangle formed by connecting with the outer corners of the eyes and the tip of the nose.

Some artifacts, such as a gap between the hairstyle and face image, may occur if the adjusted hairstyle image is merely pasted onto the face image. In our method, such artifacts are removed by deforming the hairstyle image to fit it to the target face boundary by an image-warping operation.

Finally, the hairstyle image is blended with the face image by means of an alpha-blending method, to avoid aliasing at the edges of the hair style image. The alpha value of each pixel in the hairstyle image is calculated from its intensity histogram.

Processing time depends on the size of the hairstyle image, and is about five seconds on average (on an Intel Pentium 133 MHz).

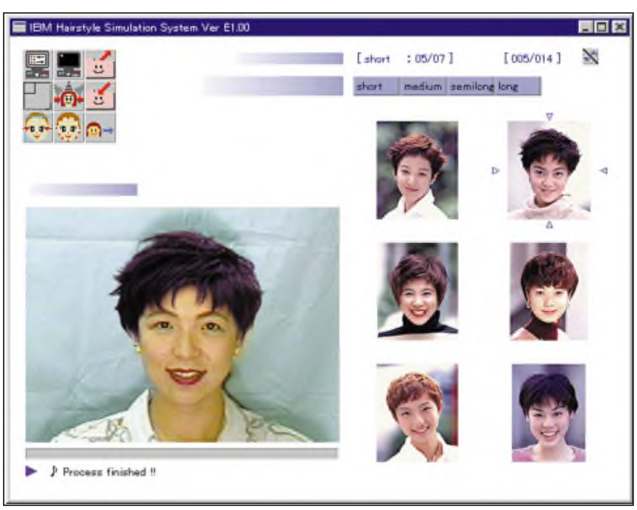

Panned/Zoomed Landscape Video Sequences Composited with Computer-Generated Still Images

Eihachiro Nakamae
Faculty of Engineering
Hiroshima Institute of Thechnology
2-1-1 Miyake, Saeki-ku
Hiroshima 731-5193 Japan
nakamae@cc.it-hiroshima.ac.jp

Xueying Qin and Guofang Jiao
Sanei Co.

Przemyslaw Rokita
Warsaw University of Technology

Katsumi Tadamura
Yamaguchi University

Yuji Usagawa
Chugoku Electric Power Co.

We present a novel, reasonable-cost approach to panned/zoomed landscape video sequence (VS) images taken by a camera on a tripod. The approach matches those images with photorealistic computer-generated (CG) still images of large-scale construction projects such as bridges and electric power transmission towers. Our newly developed compositing algorithms can be used with a consumer camcorder, a video editing system, and our own software for rendering daytime scenes under various weather conditions. Other types of equipment such as an elaborate camera controller, or special reference points in the landscapes, are unnecessary.

Features

The panned/zoomed video sequences composited with CG still images are well harmonized, in terms of both optical effects and geometrical accuracy to less than one pixel. A precise panorama composited from the CG still images and VS images is created, in which varying exposure levels due to changing camera view fields are adequately calibrated, and any geometrical distortion is eliminated.

Solving Problems

Image mosaic algorithms[1] solved the problems in making a seamless panoramic image from VS images by employing some deghosting techniques for video camera distortion due to lens/CCD. However, in our compositing of CG still images with panning/zooming VS images, much more precise geometrical matching is required. In order to composite CG still images and panned/zoomed VS images, and to make a seamless panoramic image, we solve the following problems:

1 Geometrical matching
The coordinates of the locations where the CG images should be mapped onto each VS frame are precisely estimated, and camera parameters, rotating angles, and focus length are recovered by taking account of interlaced scanning characteristics. Camera distortion is also carefully treated to prevent the CG images from floating on each VS frame.

2 Optical matching
The color and brightness between the CG images and every background VS frame, various weather conditions (the solar position, skylight, fog, haze, clouds, etc.), shadows cast by CG objects, changing of exposure due to the automatic iris of a camera, and effects of motion blur and focus blur are adequately managed.

3 Matching the different scanning images
Non-interlaced CG images are interlaced and pasted onto each VS image in order to prevent the CG images from shivering on each VS frame.

Example

Fig. a shows a panoramic image made from a horizontal panning sequence taken by an automatic iris video camera and carefully calibrated for the change of exposure. In Fig. b, a new bridge is composited with (a) by employing proposed method. In Fig. c, no calibration was considered. Fig. d shows a panoramic image made from a vertical panning sequence. Note that the new bridge castsshadows on the ground; the shadows are rerendered by using a 3D terrain model. The attached video sequences display stable effects, even when the camera is zoomed and panned simultaneously.

Reference
1 R. Szeliski, H.Y. Shum. Creating full view panoramic image mosaics and environment map. Computer Graphics Proceedings, (1997), 251-258.

Virtual Stage: An Interactive 3D Karaoke System

KeeChang Lee, ChangWhan Sul, SungJoon Hur, KwangYun Wohn
KAIST
373-1 Kusung-Dong, Yusong-Gu
Taejon 305-701 Korea
lkc@vr.kaist.ac.kr

Karaoke is a form of entertainment that enables people to become vocal stars by singing along with specially produced accompanying music. In conventional karaoke, lyrics are displayed on a television screen to help the singer follow the words. Our extended karaoke system is based on virtual reality technology. There are two motivations for this endeavor:

Firstly, Virtual Stage is a complex multimedia system that demands advanced technologies in various fields: multimedia, virtual reality, and artificial intelligence. We would like to extend the karaoke system to embrace 3D real-time graphics and animation technology, 3D interaction technology, intelligent scoring of human performance, and virtual agent technology for intelligent behavior.

Secondly, we consider Virtual Stage a common test bed for our various research activities on the following issues: motion capture/tracking of virtual agents, multi-modal interaction, methods for geometric and behavioral representation of virtual environments, and real-time management of virtual environments (rendering, simulation, and scene management).

In Virtual Stage, the input video image of the participant is shot against a blue background and chromakeyed. The keyed image is composited with the rendered image of the 3D virtual environment. The resulting scene is displayed on a large screen located in front of the participant. As a song begins, the Virtual Stage system displays the lyrics in the same way as a traditional karaoke system. But in addition to singing along to the song, the participant can interact with the virtual characters. The sensor module analyzes the keyed image of the participant, extracts necessary information such as position and posture, and generates participant-related events.

A script for Virtual Stage is associated with each song and consists of two parts: song data and interaction data. Song data, which consists of notes, lyrics, and cues, encode the underlying song to be played in the standard MIDI file format. Note data for each MIDI track are mainly comprised of MIDI note events. Lyrics data are essentially the timed MIDI meta events for each syllable. A cue is a time interval with a unique idenfier. Cues are used to separate different phases of the script such as different verses of a song, and to specify meaningful time intervals such the climax of the song. Cues are also inserted in the MIDI song as meta events. Interaction data, the other constituent of a script, controls the temporal behavior of the virtual characters. The temporal behavior of the whole script is defined by a set of rules for each interaction entity and cues defined in the song.

The virtual environment consists of objects of the following classes: avatar, agent, setting, camera, light, and music. They are called interaction entities (IEs). IEs interact with each other as the script is played. Interaction between IEs occurs according to the events generated within Virtual Stage. The events are generated from three sources: body movement of the participant, cues defined in the song, and the result of other interactions. Each IE has a set of rules, which in turn consist of the condition-action pair. The condition of a rule is a logical combination of expressions describing the basic events. As the system plays, the condition parts of the rules for all currently defined IEs are repeatedlyevaluated. When the condition expression is evaluated to true, then the corresponding action part is executed, and the rule is said to be fired.

The system consists of two computers and periphery devices such as video camera, large screen monitor, MIDI sound module, and audio. One machine equipped with the real-time video processing board is responsible for most of the processing. The video processing, the user sensing module, and graphics pipeline all run on this machine. The other machine works as the sound server. The sound server plays the underlying song by sending the MIDI messages to the MIDI sound module. Since the sound server generates the time code as well as the channel message, the progress of the song and the underlying scenario are guaranteed to be synchronized.

Acknowledgements
The work described in this sketch was partially funded by CAIR (Center for Artificial Intelligence Research) and SAIT (Samsung Advanced Institute of Techonology).

Studying Sculpture with a Digital Model: Understanding Michelangelo's Pietá of the Cathedral

Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Gabriel Taubin
IBM T.J. Watson Research Center
30 Saw Mill River Road
Yorktown Heights, New York 10598 USA
holly@watson.ibm.com

Jack Wasserman
Temple University

Digital models have clear archival value in documenting the condition of priceless works of art. Here, though, we focus on what can be learned about the composition and history of a sculpture by an expert art historian using a three-dimensional computer model rather than from direct visual inspection, traditional photography, or measurements.

Our project was initiated by art historian Jack Wasserman, who is currently preparing a book on a Michelangelo Pietá (forthcoming from Princeton University Press) now located in the Museum of the Opera del Duomo in Florence, Italy. In the course of this study, the piece has been documented in several different ways, including extensive high-quality photography and X-rays to determine the extent of damage indicated by surface cracks. Many questions about Michelangelo's composition and the statue's subsequent treatment remain unresolved.

The four figures (Magdalene, Nicodemus, Christ, and the Virgin Mary) were given different proportions in order to achieve a particular visual impact. To understand how the figures were proportioned and assembled to achieve the final effect, they are separated in the digital model and displayed to compare their relative sizes. While simple linear measurements could estimate proportions, the visual impact can be better assessed by rendering the actual pieces.

The appearance of the piece as a whole from different vantage points must also be studied. The statue was probably intended to stand further above the viewer than it does today, and it is of interest to study its proportions at different heights. By contrast, viewing the statue from directly above would allow scholars to examine the compositional flow of Christ being lowered from the cross and onto the Virgin's lap by Nicodemus.

Another issue is the impact of parts of the sculpture that were removed by Michelangelo and subsequently reattached by others. For example, when the arms of the Christ are removed, the torso of the Virgin Mary becomes visible. Altered photographs cannot show how that figure looked with Christ's arms removed, because the figure is now obscured from the front. Using the digital model, we can also reconstruct the missing left leg of the Christ, and obtain some insight on why it was removed.

On a different scale, our model will also enable a comprehensive study of chisel marks over the surface. The unfinished nature of the statue gives us a unique opportunity to understand how the artist approached his work.

Technical Approach
The size and complexity of the sculpture present many special problems. The piece stands 2.25 meters tall and rests on a one-meter base. The limbs of the figures, wrapped around one another, form a rather complex topology. Physical contact with the statue has to be minimized.

For these reasons, we selected the Virtuoso (www.visint.com) Shape Camera to scan the statue. The Virtuoso works by projecting light stripes onto the target, and then imaging the target with six side-by-side cameras. Three-dimensional points are computed using a multi-baseline stereo algorithm, and a triangle mesh is produced. A color camera mounted on the Virtuoso captures texture information.

Each surface patch obtained from the Virtuoso is about 20 cm square, and contains some 30,000 measured points. This size and resolution are design decisions built into the base Virtuoso product, which was designed for planning plastic surgery. Larger objects are obtained by merging multiple surface patches. To facilitate the alignment of patches to be merged, we projected laser dots onto the sculpture so that matching points in pairs of overlapping patches could be readily identified.

Preliminary Results
We started by CAT-scanning a reproduction of the statue sold by the museum. The scan was useful for planning the digitizing of the full-scale piece, and for preliminary renderings to show results that could be achieved.

The initial scan of the actual Pietá in Florence required over 600 individual surface patches, which were aligned using the projected laser dots. The figure shows one surface patch. We are currently developing a single triangle mesh combining all 600 patches.

Future Work
We are planning a second scan of the statue, using an improved imaging device. The new equipment will allow the capture of larger portions of geometry per pose, as well as accurate measurements of photometric characteristics of the surface.1 A position and orientation tracking system is also under study, to facilitate coregistration of multiple meshes. Two obvious areas for future work are automatic reconstruction of an accurate representation of the global geometric and photometric qualities of the statue (registration, fitting, filtering) from the collection of range images, and interactive display and manipulation of the model (multiresolution techniques, data compression, interface issues). While these problems have been studied in the past, we believe that the size and nature of the data we acquired present novel challenges.

A mesh obtained by merging three shape pictures taken with the Virtuoso Shape Camera. This particular pose produced a mesh with 23K vertices and 40K triangles. The spacing between points is about 2mm.

Session: Exploration

A Hierarchical Focus & Context Method for Image Browsing

Lars Erik Holmquist and Staffan Bjork
Viktoria Research Institute
Box 605
S-405 30 Gothenburg, Sweden
{leh,bjork}@informatics.gu.se
www.viktoria.informatics.gu.se/

Browsing collections of digital images is a common task. The collections may include photographs, scanned document pages, drawings, renderings, etc., that are located on the user's computer or on the World Wide Web. If the amount of images is large, it is important to maintain a good overview of the material,while still allowing users to examine each image in detail. Visualizations that provide such a view are sometimes called focus+context methods, since they give both an overview (context) and details (focus) at the same time.

Many applications provide overviews of image collections by presenting a set of thumbnails (smaller versions of the images) in a display area separate from the main picture. However, this is not always the most natural approach. In real life, we often put all images on the same work space, for instance a desktop. We then optimize the working space by placing the images we are currently interested in close at hand, while maintaining an overview of the whole set. When we bring an image into "focus," it does not stay in the surrounding context, since there is no thumbnail representation that exists separately from the actual image. Furthermore, images often fall naturally into categories and sub-categories (i.e., they have some kind of hierarchical ordering). This can be reflected on the work space by grouping similar images together or placing them in different stacks.

We have developed flip zooming, a focus+context approach to presenting image sets that attempts to mimic this way of working with images.[1] The main image and the thumbnails are presented on the same display area rather than in different windows. When an image is selected for viewing, the thumbnails are reduced in size and re-arranged to accommodate the main picture, which replaces the thumbnail representation. Our first implementation of the technique was a text-only Web browser, the Zoom Browser. We then generalized the method to allow browsing of any image collection.

An evaluation of our first prototype made us aware of some problems with the technique, which have been addressed in the current prototype. In particular, users found it confusing that the place where the main image would appear was not fixed; this was because we wanted to pack as much information as possible on the display. We have now made a trade-off between information density and usability, and the main image is fixed to always appear in the middle of the display.

The prototype makes it possible to group images together in hierarchies by presenting different categories in separate container areas on the display. These containers can be brought in and out of focus just like images and thumbnails. Categories can have sub-categories, which in turn can have sub-categories, allowing for any depth in the hierarchy. Since users can zoom in and out to any desired degree in each area and bring each area in and out of focus, it is possible to work deep down in an image hierarchy without losing the overview. Even though images outside the focus can become very small, it is still easy to identify categories, since the individual layout of each area is maintained.

The figures show examples of browsing a collection of images on the Web, in this case portions of the collection of The Louvre Museum (www.louvre.fr). The categories are Prints and Drawings; Paintings (with sub-categories for the French, Italian, Flemish/Dutch/German, and Spanish/English schools); Greek, Etruscan, and Roman Antiques; and Egyptian Antiques (with sub-categories for sculptures and paintings). For browsing such image collections, we believe our method can present a useful alternative to traditional image-browsing applications for local file systems and the Web.

Reference
1 Holmquist, L.E. Focus+context visualization with flip zooming and the zoom browser. In Extended Abstracts of CHI '97, ACM Press 1997.

Time Travels in Virtual Landscapes

Ingo Braun, Anja Kutzner
Kulturbox GmbH
ingo@kulturbox.de, anja@kulturbox.de

Birgit Böhme, Uta Simmons
Centre for Berlin Studies
birgit@kulturbox.de, uta@kulturbox.de

Askan Striepe, Christian Quintus
Fraunhofer Institute for Production Systems and Design Technology
askan.striepe@ipk.fhg.de
fquintus@e27.com

Andreas Knoche, Leonie Schäfer
Technical University of Berlin
knoche@cs.tu-berlin.de
leonie@acm.org

Imagine wandering through tBerlin in the 19th century, discovering old streets and places, having a coffee break at a traditional cafe, reading the newspaper of the day, and later on having a look into the windows of a shop nearby. While exploring the city, you come across people dressed in contemporary clothes, and you see objects from former times in their historical environment.

What did Berlin look like? What was the atmosphere? What was the talk of the day?

This sketch presents a project that aims at development of a four-dimensional online information system that enables the user to find answers to these questions. The system is four-dimensional, because users can not only explore a three-dimensional cityscape, but they can also observe its changes over the course of two centuries. The main emphasis of the system is the presentation of historical information in an interactive and immersive way, so users will be encouraged to discover history rather than merely gather and consume information.

The cityscape shown is a digital reconstruction of Friedrichswerder, a central quarter of the city of Berlin. The scenery, its buildings, and objects are realized in VRML97. The architectural models were reconstructed with the use of a CAAD program and then converted to VRML97. Additional scene objects like vehicles, goods in window displays, or advertising posters are modeled in VRML with textures generated from old illustrations or photographs of the original historical objects. The time-traveling functionality and the city map navigation tool are realized in Java.

Objects exhibited in the scenery were made available by museums and libraries cooperating with the project. Contemporary radio and television broadcasts are provided by radio archives. A database connected with the virtual scenery contains further detailed historical information on each object, image, or broadcast shown. This database not only offers the opportunity to search for additional material not displayed in the virtual landscape. It also allows the user to follow a certain string of cultural history (for example, media history).

Acknowledgements
This work was financed by DFN e.V. with funds from Deutsche Telekom Berkom GmbH. Architectural data were provided by BAUAKADEMIE Gesellschaft für Forschung, Bildung und Entwicklung mbH.

Session: Exploration

Smithsonian Without Walls

Thomas Mueller, Marcelo Ronchini
razorfish inc.
324 Pearl Street #3A
New York, New York 10038 USA
thomas@razorfish.com

When the Smithsonian Institution decided to create an innovative digital exhibit to extend its vast collection of artifacts and curiosities beyond the physical confines of a building, it partnered with Razorfish and Plumb Design to further develop the concept, determine functionality, and create the working prototype of Smithsonian Without Walls – Revealing Things.

The exhibit shows visitors the unique relationships between everyday objects and allows them to learn and understand the important cultural and personal history of the various items. The prototype's navigation is integrated with content using Plumb Design's Thinkmap, a Java-based application that creates interactive displays that allow users to understand the relationships among complex data. Visitors to the online exhibit can easily explore the multi-dimensional relationships among objects by era, theme, or the objects themselves. The interface design eliminated the traditional role of navigation as a supplement to content by fully integrating it into the xploration process.

Pictures, video, music, and spoken word, streamed via Macromedia Flash, communicate stories and memories as seen through the eyes of the people who make, use, and study the objects. These narratives and accounts allow users to comprehend what makes these objects special and gain increased awareness and understanding of their own special surroundings.

The fully functioning exhibition will provide users the option to save their pathway through the site and share it with others. Users can also enhance the exhibit by donating their own objects.

Features

- Unique Java-based interface and navigation for access to online archives
- Games and puzzles
- Comment book
- Object locator
- Online tours

Benefits

- Expands the reach of the Smithsonian Institution by increasing availability of the exhibition
- Solidifies the leadership role of the Smithsonian Institution in preserving cultural heritage

Applying Depth-of-Field Effects to Power Substation Simulation System Using Virtual Reality Technique

Koichi Arai and Shigenobu Furuta
The Kansai Electric Power Co.,Inc
3-11-20, Nakoji, Amagasaki
Hyogo, 661-0974 Japan
arai@rdd.kepco.co.jp
sfuruta@rdd.kepco.co.jp

Arata Watanabe and Noboru Kamizi
Toyo Information Systems Co.,Ltd.
9-1, Toyotsu-cho, Suita
Osaka, 564-0051 Japan
arata@karl.tis.co.jp
kamizi@hp99.tis.co.jp

This indoor distribution substation simulation system was developed for use in training of substation maintenance personnel. The simulator allows users to go through a routine inspection procedure involving a variety of operations as they "walk" through a virtual substation. Its principal functions include visual and aural simulation of situations that maintenance personnel may encounter in an actual substation. But some people who experienced the simulation felt that the operation in virtual space was difficult to achieve. The problem came mainly from difficulty in judging the distance from the eye-point to the object to be reached in virtual space. An experimental study was conducted to apply depth-of-field effects techniques to the simulation system in an attempt to improve the sense of distance.

Introduction
In visual presentations where depth-of-field effects are not applied, users, whether looking near or far, can see all objects clearly. But it is only natural that users get blurred images of the objects unless they are staring at them. In this study, the multipass rendering technique is used to examine the depth-of-field effects on moving pictures. Some OpenGL routines were added to the IRIS Performer and applied to the simulation system.The simulation system is based on IRIS Performer2.1 on an SGI Onyx RE2.

Generation of Depth-of-Field Effects
The method to generate depth-of-field effects is as follows:

1 Renewal of viewpoints
2 Culling (for each viewpoint)
3 Clear accumulation buffer
4 Start jittering loop
5 Set the view in response to jittering
6 Clear frame buffer
7 Draw scenes and output them to frame buffer
8 Divide content of frame buffer by the number of jitterings and add them to accumulation buffer
9 Copy content of accumulation buffer onto frame buffer and output them to display

The entire process is generated in the simulator's limitless loop, and steps 4-8 are in the jittering loop. Steps 3, 5, 8, and 9 are realized by using OpenGL. And steps 6 and 7 are realized by using IRIS Performer.

Results and Conclusion
The number of jitterings was changes experimentally. The frame rate varied depending on the number of jitterings. Consideration must be given to a trade-off between the visual effects and the response time. Searching experimentally for the adequate number of jitterings for gaining the desired depth-of-field effects on a real-time simulator is required. As for the frame rate, keeping more than 10 frames/second is absolutely necessary. Since the frame rate depends on a workstation's capability, better workstations should be used to generate better frame rates.

It is feasible to apply the technique of depth-of-field effects for real-time simulators. Parameters for depth-of-field effects were defined and generated on Onyx RE2. Multipass rendering has a promising application for functional enhancement of IRIS Performer.

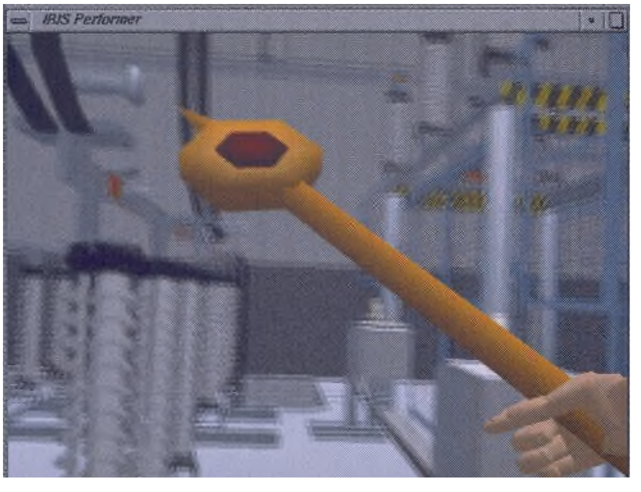

A scene in the substation with depth-of-field effects (the user is holding a voltage detector).

Coupled Models for Visualizing Respiratory Mechanics

Jonathan M. Kaye, Dimitris N. Metaxas, Frank P. Primiano, Jr.
University of Pennsylvania
200 South 33rd Street
Philadelphia, Pennsylvania 19104-6389 USA
kaye@linc.cis.upenn.edu
www.cis.upenn.edu/~kaye/research.html

Recently, there has been an explosive growth of multimedia software designed to enliven presentation of medical topics through computer animations, actual videos of patients, and quantitative patient models. However, we have yet to see programs that combine sophisticated physiological models, such as those developed in the bioengineering literature, with 3D anatomy in interactive simulations. We see such a combination as the basis for a learning tool in which the student is given the greatest freedom to explore consequences of disease and intervention.

We have created a 3D simulation environment for modeling respiratory mechanics and certain cardiopulmonary interactions.[1] Because of the domain and our intended application, we can drive a high-dimensional system, namely the 3D anatomy (the "anatomical modeling") by one with fewer parameters: a compact, scalar description we call the "physiological modeling." We created the 3D anatomy by reconstructing organ geometry from CT scans. The 3D anatomical modeling, involving kinematics and deformable body dynamics, gives a qualitative perspective on the anatomical changes. The physiological modeling is expressed in quantitative terms familiar to the student and clinician.

Our vision is to produce compelling, accurate, interactive environments for studying normal and pathological behaviors in cardiopulmonary physiology. This can be beneficial for both education and clinical practice in which the visual feedback and the predictive capabilities of the system may be helpful in simulating situations that would be difficult to control clinically.

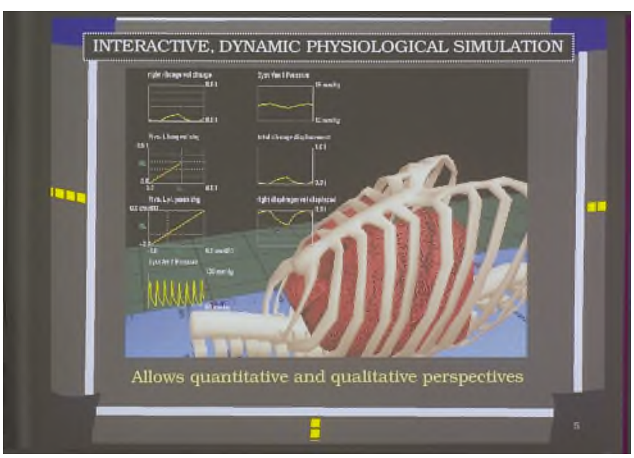

3D Anatomical Modeling

Our 3D organ models deform during the simulation using a combination of kinematics and dynamics. Motion is achieved for the rib cage, torso, and parietal pleura (the membrane enclosing each lung) using kinematics by interpolating 3D geometry according to displacements calculated in the physiological modeling module. The lungs are controlled by a combination of kinematics and dynamics, as dictated by the transmural pressure gradient (the difference of the pressure inside the lung and outside) across the lung surface. To simulate the lung's anisotropic shape change for large displacements, we interpolate a "core" shape based on the pressure gradient. The core is attached via zero rest-length springs to a surface mesh. The surface mesh is a spring-mass system responsive to both internal forces and contact forces derived from the transmural pressure gradient, applied along surface normals. The 3D mechanical parameters are estimated from the organ reconstruction models. For visual fidelity, we keep the lung from penetrating the parietal pleura by canceling forces on the lung perpendicular to the surface of the parietal pleura.

Physiological Modeling

The physiological models consist of mathematical descriptions (expressed as ordinary differential equations) of the mechanical relationships between components. The model parameters, such as resistances and compliances, are readily understood and clinically meaningful. The model variables are changes in pressure, volume, and flow. Although the cardiovascular model is typical of those appearing in the bioengineering literature, our respiratory model is novel because of its multi-compartment chest wall.

Reference

1 Kaye, J., F. P. Primiano, Jr. and D. Metaxas. A 3D virtual environment for modeling mechanical cardiopulmonary interactions. To appear in Medical Image Analysis, 2(2), 1998.

Language Through Gesture in a VRML World

Sarah Geitz, Chae-hi Park, Diane Brentari
Purdue University
West Lafayette, Indiana 47906 USA
sgeitz@tech.purdue.edu

Stephen Maher
NASA/Goddard Space Flight Center
Greenbelt, Maryland USA
maher@holodeck.gsfc.nasa.govs

Creating interactive, three-dimensional worlds for view in the essentially two-dimensional, window-like environment presented by a computer monitor requires making use of a number of cognitive cues to convey essential information. Additionally, designing a VRML world necessitates economy in file size to facilitate downlaod speed . Maintaining a small file size requires isolating and using only the essential information needed to conveys pecific meanings. In order to isolate the essential visual cues necessary to communicate meaning through gesture, we must first examine how the context of the viewing environment influences how the information is perceived by the viewer. The results of human testing of two avatars designed to teach ASL in a VRML world demonstrate how context influences the essential information required to communicate language through gesture.

Problem Statement

Virtual Reality Modeling Language (VRML) allows for transmission of animated, user-controlled, three-dimensional graphical objects and scenes over the WWW. VRML enables one to change one's point of view at will in an intuitive, natural way that can potentially simulate human face-to-face interaction better than more passive technologies such as video. A scene or model can include objects, colors, texture maps (photographic images that are overlaid onto objects), and camera locations. The advent of this technology means that, for the first time, animated models and environments of an interactive 3D nature can be transmitted to geographically and educationally diverse audiences. Specifically, our project investigates the effectiveness of using animated 3D models in teaching a visual language.

All natural languages recursively combine phonological primitives to create words and utterances. Instead of features such as voice, coronal, lateral, etc., the phonological structure of ASL recursively combines features of movement, hand shape, orientation, and place of articulation to create words. The capacity to visually comprehend movements from varying points of view and effectively use the space around one's body is critical to the learning and communication process. Sign movements are inherently 3D, yet until recently the means to record movement in a 3D venue have not been available. Visual documentation of sign languages consists of orthographic drawings, photographs, films , videotapes, and CD-ROMs. The use of static representation to render the shape of movements, such as drawings or photographs, is not effective, especially in teaching. Moreover, phonetic and phonological notation systems, such as Hamnosys (Prillwitz, 1993), are useful only to those individuals who have spent considerable time learning them, in addition to learning to sign. Video is one alternative to static pictures as a teaching aid, but video places the viewer and the signer at a fixed angle to one another, one that may enhance perception of certain aspects of the signed utterances and minimize others. Furthermore, since signingi nterlocutors may not always be standing directly face-to-face, videos do not help the student see signs from a side angle or when looking upwards or downwards towards a signer who is of a different height. Thus, traditional means for recording sign languages do not adequately represent the 3D aspects of the language in a manner conducive to intuitive understanding of a novice signer.

A virtual environment provides a spatial representation of an object, system, or performer. The theoretical advantage to using a spatial form of representation such as VRML as an auxilary teaching and recording tool for ASL is that the virtual environment provides for physically close, repeatable, and lengthy observation of the virtual performer. However, the type of close observation provided within a VRML world would be socially unacceptable in the real world. Direct, intimate manipulation of the signer within virtual space allows the viewer to perceive hand shapes and gestures in a manner they could not perceive them in the real world.

Procedure

The challenge is to design 3D models in a VRML setting that provide a non-linear, spatial vehicle for representing kinetic events in a way that allows their use by human cognitive mechanisms. Determining the level of graphical detail required to create an ideal representation of the humanoid for the purpose of communicating gestural information within VRML requires testing the models on human subjects. We examined the effectiveness, first, of a static virtual model developed to augment the teaching of finger spelling. We then modified this avatar based on our initial findings and retested the model presenting the same information, plus several animated signs. All testing was conducted using subjects unfamiliar with ASL.

Findings

When a model designed to represent a human hand with finger digits of varying lengths was used, subjects exaggerated inconsistencies produced by differences in finger length. Because the interactive nature of the VRML world allowed for closer inspection of the hand, we found the VRML environment necessitated modifying the lengths of digits of the fingers of the avatar, thus producing a cartoon like figure.

Significance of the Investigation

Evaluation of the differences in sign reproduction when 2D and 3D models are employed pedagogically can help not only in understanding how best to model the figure of an idealized virtual signer, but also in identifying differences in human perception of cognitive cues relating to communication through gesture in VRML worlds in general. It can help in our understanding of how to design avatars that better communicate the transmission of human presence.

Reference
Prillwitz, Siegmund. Hamburg notation system for sign languages: An introductory guide. Hamburg, Germany: Signum Press.,1989.

Internal Representation by the Magic Light

Seiki Inoue
ATR Media Integration & Communications Research Laboratories
2-2 Hikaridai, Seika-cho, Soraku-gun
Kyoto 619-0288 Japan
sinoue@mic.atr.co.jp

There are many systems that can handle three-dimensional objects in real-time virtual worlds. These systems allow an object to be perceived from various directions and allow observers to feel as if they were actually holding and examining the object in their own hands. This is the most common way to appreciate objects. Here, the authors break with convention and propose a new method for representing objects.

Figures show one example of creating virtual images that is impossible in reality. When the observer moves the spotlight on the displayed object, the portion of the object illuminated by the spotlight becomes transparent, and what is hidden inside the object appears. The observer can experience a sense of discovery.

Three processes are necessary to realize this representation:

1 Determining the location and direction of the light.
2 Detecting the region illuminated by the light.
3 Adjusting the degree of transparency of the illuminated region.

These processes all need to be conducted in real time. The first process is made possible by utilizing three-dimensional sensors that are commonly used in virtual reality. For the second and third processes, CG techniques are necessary. Although rea-time CG techniques are progressing rapidly, it is difficult to flexibly change processing because hardware that performs high-speed calculations must be changed.

It is easy to change normal light sources or adjust the brightness of the object according to the movement of the light source with present hardware. On the other hand, present hardware cannot calculate the illuminated region by the light source and alter the degree of transparency within that region. For that reason, software processing is necessary, but it cannot be performed in realtime. The authors have devised a new method by combining three-dimensional CG techniques with image synthesis techniques. We use CG techniques as the key for synthesis instead of the output image. The external and internal images of a three-dimensional object are made and recorded on two VCR units. The CG workstation outputs the key information for synthesis. This key image is the illuminated region of a textureless object and can be easily produced with normal three-dimensional CG. It corresponds to the change of the light source location in realtime.

The three-dimensional location sensor is comprised of two units, one for sending and the other for receiving. The receiving unit is placed under the desk so that it is not in view of the observer.

Significantly, the system utilizes the image generated by CG as the key information. By utilizing this key information, it is possible to create the image with the illuminated region appearing as if it were transparent. This new representation method gives excitement, enjoyment, and discovery to users. It can also provide creators with a powerful tool for making interactive CG content. Creators can embed many items in a scene, and users can discover them with joy. Several textures can be mapped on an object, and several objects such as a Russian Matryoshka doll are included.

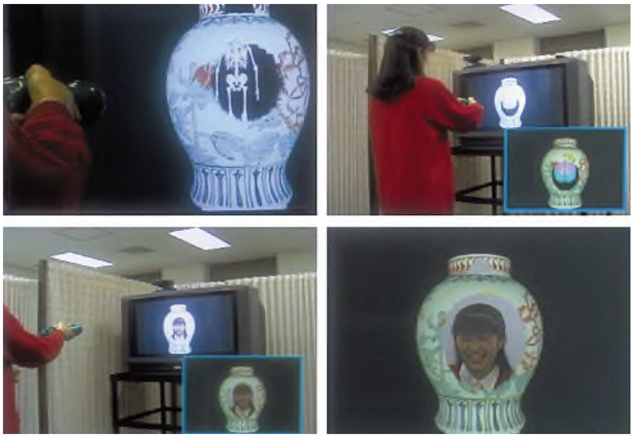

The Making of Black-Hole and Nebula Clouds for the Motion Picture "Sphere" with Volumetric Rendering and the F-Rep of Solids

Gokhan Kisacikoglu
Cinesite Digital Studios
1017 North Las Palmas
Los Angeles, California 90038 USA
kisa@cinesite.com

The implementation consists of the volume renderer itself and the implicit shape module (ISM).

Ray tracing through a volume of gas computed the sum of the densities and density-weighted colors, called samples. The contribution of the samples attenuates with the cumulative density as the ray advances away from the camera. The ray tracing proceeds until the opacity clips or the ray leaves the complete volume.

The volume rendering approach taken for this project is computationally intensive. The early-version volume renderer developed at Cinesite supports point-like and directional light sources, with full self shadowing and asymmetric scattering. Ray tracing consists of traversal of voxels and cells, accumulating radiance and attenuation. For the purposes of this project however, the gas was lit internally, so it acted as a glowing and attenuating material with a variety of colors and complex structures. To accomodate the artistic design process and be as flexible as possible, the renderer accesses material properties through method calls in an API. This allows artists and technical directors to build volumetric data in procedural codes without dealing directly with the rendering process.

The ISM mainly provides the point samples as the ray tracing takes place, and it manages the primitives, sampling and interpolation of these samples, and loading and evaluation of the appropriate shaders for these primitives.

The volumetric shaders used to represent nebula-like structures require varying densities and colors. Noise patterns were used in both shading and displacing the isosurfaces. Briefly, the fractals and random patterns are associated with the f-rep of the solids when inside of the clouds (defined as $f(x, y, z, t) >= 0$) and another atmospheric (ambient) shading when outside of the shape (defined as $f(x, y, z, t) < 0$). The properties of implicit surfaces were also used to enhance the known shading techniques.

We have optimized the f-rep evaluation. The cloudy shapes we needed in order to achieve the desired nebula effects required large databases. A "realistic" look for these gas clouds requires certain shapes based on storyboards from pre-production, as well as examples from real counterparts. Independent of the choice of the implicit functions, the algorithm to solve the f-rep and evaluate the solids was designed to be capable of localizing the data in a smaller volume to reduce the volume of operations.

Executing the shader for the whole ray trace was very expensive, and in some cases such as evaluation of gas motion in vectoral fields, it was impossible directly. So we investigated methods to sample and store volumetric data efficiently with fast access and interpolation for a specific point throughout the ray trace. The resolution increased toward the edges of the clouds, as the most detailed information is rendered there.

All the modules were written in C++. The volume rendering software was developed by Cinesite Digital Studios. The ISM and shading libraries were developed during production. The volumetric renders were conducted on our rendering servers.

Volumetric rendering by using f-rep of solids enabled us to design and create gas clouds and complete the project on time. Given our drawing and compositing skills, the resulting animations would be much more difficult to create by other methods.

Credits
Special Visual Effects: Cinesite Digital Studios
Visual Effects Supervisor: Carlos Arguello
Visual Effects Producer: Aaron Dem
Associate Producer: Chris Del Conte
Flight Deck Hologram Sequence
CGI Supervisor: Fernando Benitez
CGI Technical Director/Author: Gokhan Kisacikoglu
CGI Animators: Cesar Velazquez, Richard Klein
R&D Supervisor: Tom Asbury
R&D Programmers: Jerry Tessendorf, Scot Shinderman, Mayur Patel
Executive Producer: Gil Gagnon
Production Manager: Kristen Niederholzer
CGI Department Supervisor: Joe Pasquale
Production Coordinators: Cathy Cassese
Production Assistants: Tony Sands, Christina DeSilva
Animator Assistants Supervisor: Vincent Lavares
Animator Assistants: Ariel Y. Oclarino, Jackson Yu, Brian Shows
Lead Data Engineer: James Blevins
Supervisor: Tony Sgueglia
Supervisor: Brian Bud Kun
Data Engineers: Pavel Dvorak, M. Ruth Vasquez
Visual Effects Editor: Rod Basham
Assistant Editor: Steve Rhee
Avid Editors: John Alvord, Amnon David
Art Director: Nicolai Strehl
Conceptual art department: Don Campbell, Jerald Stuart
Director of Technology: Guy Griffits
Engineering services: Mark Girard, Rik Bomberger, Derek Laing

Web-Based Sonification of Space Science Data

Robert M. Candey, Ramona L. Kessel
NASA Goddard Space Flight Center
Code 632, Bldg 26 Room 126
Greenbelt, Maryland 20770 USA
Robert.M.Candey@gsfc.nasa.gov
cdaweb.gsfc.nasa.gov/

Justin R. Plue
Eleanor Roosevelt High School
7601 Hanover Parkway
Greenbelt, Maryland 20770 USA
plue@nssdc.gsfc.nasa.gov

Sonification is the use of non-speech audio to convey information[1] (with sonar and Geiger counters, for example). The best example in space physics was the use of sound for detecting micro-meteoroids that impaced Voyager 2 as it traversed Saturn's rings; these impacts were obscured in the plotted data but were clearly evident as hailstorm sounds.[2]

We are extending the Coordinated Data Analysis Web (CDAWeb) data browsing system (cdaweb.gsfc.nasa.gov/) to display space science data as sounds in addition to its existing plot capability (line plots, spectrograms, images, etc.). In CDAWeb, the user selects spacecraft, instruments, datasets, and time periods of interest and is then presented with a list of variables to display. PERL and Interactive Data Language (IDL) routines on the CDAWeb server interactively create the selection pages and then read the data from Common Data Format (CDF) files (used for its portability and self-describing meta-information), plot the data (automatically selecting reasonable plot types and scaling), and return the plots as GIF images.

We are extending CDAWeb to provide an alternative display type for certain variables where an IDL routine writes the data out in a JavaScript routine that calls the Beatnik (www.headspace.com/) plug-in to the user's Web browser. When the user clicks on the sound data icon, the Beatnik plug-in plays the data array as varying pitches, loudnesses, or drumbeat rhythms.

Sonification Background

Sonification is useful for monitoring some data while looking at something else, for complex or rapidly/temporally changing visualizations, for data exploration of large datasets (particularly multi-dimensional datasets), and for exploring datasets in frequency rather than spatial dimensions. A good reference to research on the use of audio representation of data is the International Conferences on Auditory Display (www.santafe.edu/~icad/).[3]

Complex datasets (e.g., particle measurements varying in energy, look direction, time, and particle species such as electrons and ions) are usually only examined in a subset of dimensions at a time, forcing researchers to build up a picture in their minds of the whole dataset. Sound can be used to represent these other dimensions, using pitch, loudness, rhythm, damping or attack/decay rate, direction, duration and repetition, timbre and harmonics, phase, and rest periods.[4] Audio can also be used to reinforce the visual displays as alternative ways of looking at or hearing the data, particularly where visual patterns are hidden until identified by other means. An additional benefit of sound for visualization is that the ear is more sensitive to different frequency bands and patterns than the eye. Also, auditory accompaniment (as shown for movie music) clearly leads to improved visual information.

Besides improving data exploration and analysis for most researchers, the use of sound is especially valuable as an assistive technology for visually impaired people and can make science and math more exciting for high school and college students.

Future

The sonification technique shown here is very rudimentary. We hope to improve this tool based on others' research in areas such as psychoacoustics, semiotics, ergonomics, commercial sound design, and music. In particular, we plan to call the JavaSound API (which incorporates Beatnik in JDK 1.2 (www.javasoft.com/) as part of the extension of CDAWeb to use Java for more user control and interactivity, including the option to use other dimensions of sound. Eventually, we also want to try representing plasma-wave data with sound loudness representing intensity, pitch for frequency (or energy in particle data), and left-right amplitude and phase for polarization differences. The user would mouse over the plot of multispectral data and play sounds correlated with some of the data parameters, similar to a doctor listening with a stethoscope (as proposed by Wolff[5] and others).

Conclusion

Sonification will add to our scientific research capabilities, especially in complex multi-dimension and multi-dataset research such as for Earth and space sciences. It will also complement existing visual displays, and it may identify new phenomena that current display techniques miss.

References

1 International Conference on Auditory Display (ICAD) members. Sonification: A status report and research agenda,. White paper to National Science Foundation, 1998.
2 F. L. Scarf, et al. Voyager 2 plasma wave observations at Saturn. Science, 215, 29 January 1982, 587-594.
3 Gregory Kramer. Auditory display: Sonification, audification, and auditory interfaces. Addison-Wesley, Reading, Massachusetts, 1994.
4 E. S. Yeung, Pattern recognition by audio representation of multivariate analytical data. Anal. Chem., 52, 7, 1980, 1120-1123.
5 R. S. Wolff. Sounding out images. Computers in Physics, 6, 3, 1992, 287-289.

Lagrangian Visualization of Natural Convection Mixing Flows

Luis M. de la Cruz, Victor Godoy, Eduardo Ramos
Universidad Nacional Autonoma de Mexico
Circuito Exterior de C.U., Coyoacan, C.P. 04510
Mexico D.F., Mexico
lmd@labvis.unam.mx
victor@labvis.unam.mx
erm@ds5000.super.unam.mx

Deformable media motion visualization has been used for centuries to extract information from dynamic flows. For instance, we know that Leonardo da Vinci used to throw pollen grains to follow the vortices on water streams and thus visualize geometric forms that he then painted in his brilliant sketches. In our day, a number of experimental techniques for flow visualization take advantage of the optical, or other physical, properties of the working fluids. On the theoretical side, numerical modeling of fluid flow has advanced noticeably in recent years, thanks to the increase in processing speed and memory capacity of digital computers. These new tools have been particularly useful in time-dependent flow studies where a high spatial resolution is required.

It is now feasible to model turbulent flows with low Reynolds numbers where spatial resolutions of $(1/300)^3$ of total volume are required. The result of these models is a set of three-dimensional fields of velocity, pressure, and temperature whose size depends on the grid fineness. In transient flows, at each time step we have velocity and temperature fields that describe the dynamics and heat transfer of the flow. Once we have detailed space and time information about the flow, we need tools to extract this information in a simple, understandable way. Due to frequently complex three-dimensional topology, one-dimensional or bi-dimensional sub-sets of fields do not clerly reflect some of its properties. For this reason, it has been necessary to devise new techniques to visualize which show flow characteristics of interest.

These techniques include isosurfaces, particle tracking, definition of secondary variables, and their display (vorticity, heat, flow, etc.). Also. it is possible to construct graphs with a high degree of realism by using tools like shadowing, perspective, texture, and especially animation. These computational tool schemes can be of enormous use in flow analysis in natural convection with mixing, where fluid convolutions can be extremely complex. In this work, we present an example of this flow using Lagrangian tracking and sophisticated graphic techniques.

Theoretical Model

We consider a cubic prism filled with a Newtonian fluid and with Prandtl number of 200. The gravity is pointing in the negative vertical direction. The left and right walls of the prism are kept at a constant temperature in all areas but oscillating in time. The rest of the walls are considered adiabatic. Flow details are obtained by numerical integration of time-dependent governing equations of mass, momentum, and energy.

Visualization Technique

The result of solving the equations of balance is a set of pressure, temperature, and velocity fields as time and position functions. However, contrary to popular perception the "unprocessed" Eulerian velocity gives little information about mixing, and the typical ways of visualizing a flow (streamlines, pathlines, and to a lessen degree, streaklines) are insufficient to completely understand the process.

Lagrangian Tracking of Surfaces

Lagrangian particle tracking consists of determining the geometric position of points of space that were occupied by a "particle" of fluid at a given time. The procedure to do this comprises two steps that are iteratively repeated. The first step consists of defining a point in the integration volume. This is the initial point $\vec{x}_i$ where a particle is at initial time t_i. In the second step, we calculate the final point x_f, where the particle is after interval $\delta t = t_f - t_i$, with the formula: $\vec{x}_f = \vec{x}_i + \vec{U}(\vec{x}_i, t_i)\delta t$. The velocity $\vec{U}(\vec{x}_i, t_i)$ is evaluated at the initial point at the initial time. Once the value x_f is known, this point is redefined as initial, and the process is repeated.

Langranian tracking of surfaces is the direct generalization of particle tracking. We define the initial surface inside a fluid as a set of points x_i localized in convenient but arbitrary positions. At time δt, the surface is defined by a set of points x_f. The deformation of the surface will give information about the general dynamics of the flow, and in particular it shows zones where neighbor points will separate due to local stretching, or zones of the volume where the surface folds, permitting points that were originally far away to come together.

In this work, two spherical surfaces centered in the cubic prism and with radii of 0.4 and 0.2 respectively are defined. The result of Lagrangian tracking is composed of a set of data containing the position of two spheres at subsequent times. The numerical solution of balance equations was carried out in an Origin 2000 computer at DGSCA-UNAM. The characteristic duration of a run required to calculate an adimensional time unit was 1,270 segs. Visualization of the deformed surface was created in a Silicon Graphics Onyx computer with six R10000 processors. The characteristic time to generate an image using raytracing is approximately 240 segs of CPU.

Examples of oscillatory natural convection flows are presented. In all cases we considered Ra = 5×10^5 and four hot-cold cycles. Shadow tools, texture mappings, clipping planes, and color texture animation were employed to give more realism to surfaces. See the figure.

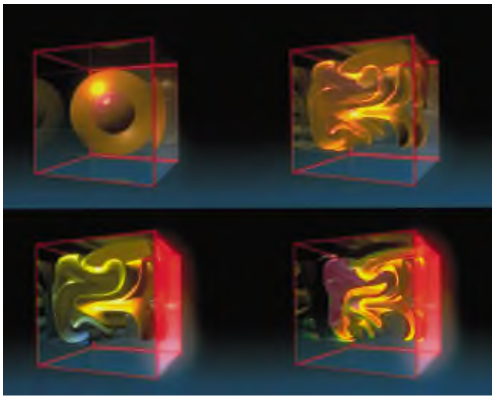

TVML (TV Program Making Language)

Masaki Hayashi
Multimedia Services Research Division
NHK Science and Technical Research Laboratories
1-10-11, Kinuta, Setagaya-ku,
Tokyo 157 Japan,
hayashi@strl.nhk.or.jp

The method used for producing TV programs has been improved over the long history of television and is now established as a powerful way of presenting information. However, TV programs are still only produced as professional video productions, and it is still hard for non-professoinals to produce complete TV programs.

We have developed software that can produce full TV programs in real time by using rea-time computer graphic (CG) characters, a voice synthesizer, and multimedia computing techniques. With this system, users can create their own TV programs on a desktop workspace simply by making a text-based script written with TVML (TV Program Making Language), a description language that we have designed for this purpose.

Automatic TV Program Generation by TVML

In order to allow a computer to understand a script written by a human and convert it into a TV program, we have developed a language that can be understood both by human and computer. TVML is a text-based language that can describe all the necessary elements for a TV program production. We designed the language to have high-level abstract expressions (for example "display title#1," "camera zoom in," and so on). On the other hand, TVML Player is a software system that can interpret the script written in TVML and generate audio and video of the program in real time. A platform for TVML Player consists of a graphic workstation and a voice synthesizer. It produces a program by combining TVML scripts and various data (audio, video, and computer graphics). TVML covers: studio shot susing real-time CG characters with synthesized voices, CG studio sets and CG camerawork; movie file playing; titles using HTML-like text layout language; superimposing; sound files played as background music; and narration by synthesized voices.

What Kind of Language is TVML?

In a TVML script, one line corresponds to one event. TVML player performs one event, and when the event is completed, the next event described on the next line is performed. Here is an example of TVML event description:

If you want to make a character Bob say "My name is Bob," in a studio shot, it's written as:

character: talk(name = BOB, text = "My name is BOB.")

The "character" segment is called event. The "talk" segment is called command prepared for the event. In parenthesis are the various parameters used for the command. There are 11 types of events: character, camera, set, prop, light, movie, title, sound, super, narration, and video. A user can make TV programs by writing out the events described above one after another.

Applications of TVML

We have developed TVML Player Ver1.1 based on a language specification of TVML Ver1.0. It is distributed as freeware. Currently, the hardware platform required for the TVML Player consists of a graphics workstation (Silicon Graphics, Inc.) and a voice synthesizer (Dectalk Express).

Applications of TVML might include:

1 A TV program production tool for non-professionals.
2 A professional tool for producing simple programs such as event guides, educational programs, etc.
3 A training tool and simulation tool for TV directors
4 TV program generation at the viewer's site by implementing the TVML player in a TV set.

DView

Patti Koenig
Jet Propulsion Laboratory
5849 Buena Vista Terrace
Los Angeles, California 90042 USA
koenig@helios.jpl.nasa.gov
dview.jpl.nasa.gov

At the Jet Propulsion Laboratory, there is a need for fast 3D visualization, rapid prototyping, and real-time analysis of articulated spacecraft, their trajectories, and instrument field of views. From as early as the 1960s, Jim Blinn and others have produced software to meet the needs of JPL mission visualization. Dview, a new language and accompanying software, is the result of one such effort. There are a number of 3D modeling packages available. However, Dview is distinctly unique because its language is based entirely on construction of simple command-text strings, and absolutely no programming protocol is required. A text string syntax is useful because it avoids problems that are typical of other 3D graphics software.

Model construction and animation involves significant programming in traditional languages. For example, knowledge of Java programming is required for Java3D, and C++ programming is required for Open Inventor. The goal of Dview is to provide the same functionality as these software packages but reduce the amount of programming. Because the entire Dview language is based on a simple text string, Dview can "plug" in to any programming language (C, C++, Fortran, TCL, Java, or even a custom language). The interface to Dview is as simple as a single subroutine or socket connection. This makes sophisticated 3D graphics instantly available to any application. Users have the option of using Dview's user interface or creating their own interface to the Dview graphics engine.

A common difficulty with modeling software is that the user is tied to a proprietary environment. For example, the 3D graphics modeler Alice (alice.org) and Geomview (www.geom.umn.edu/software/geomview) provide some functionality similar to Dview, but they both have their own unique programming syntax and file formats. Dview bypasses these restrictions by allowing its text commands to be incorporated directly into any language, so no extraneous knowledge is required except for the simple Dview command text-string syntax. There might be a natural tendency to think that Dview allows for only simple graphics applications; on the contrary, Dview is capable of state-of-the-art computer graphics with kinematics. It supports all the geometry primitives of VRML as well as additional parameterized primitives that are exported to VRML as indexed face/line sets. Dview animation commands have an intuitive syntax and are convenient for automatic generation of VRML animation interpolator nodes.

Models constructed using Dview commands can export directly to VRML 2.0, and animations using Dview commands can export directly as VRML 2.0 animations. Dview can also save animations as a series of numerically sequenced GIF or RGB files for creation of animated GIFs, MPEG files, or QuickTime movies. Dview has a complete GUI with its own model builder that allows quick interactive construction of articulated models.

Dview was implemented using X/Motif & OpenGL. A second compatible implementation using Java AWT and Java3D is underway. Dview executables, instructions, and extensive examples can be downloaded via a free click-wrap license agreement from the JPL Web site (dview.jpl.nasa.gov).

Session: Production

Previsualization for "Starship Troopers:" Managing Complexity in Motion Control

Colin Green
Pixel Liberation Front
1316 1/2 Abbot Kinney Boulevard
Venice, California 90291 USA
colin@thefront.com

Many film productions find it difficult to effectively manage production of extremely complex visual effects. In producing these scenes, directors and cinematographers often lack the creative control they would prefer to have, especially when the scenes involve numerous models and extensive motion-control photography. Complex visual effects have also proven to be hard to plan, both in terms of their cost and in the time it takes to execute them. 3D animation can aid in this process. Although exchanging data between computers and motion control systems is a fairly well-established practice, the process only works well with relatively uncomplicated shoots involving, for example, a single set-up and a single element. It tends to break down in more complex situations involving numerous rigs, camera systems, and photographic elements.

For the recent film "Starship Troopers," Pixel Liberation Front, working with Sony Pictures Imageworks, developed a new process that enabled it to create and manage motion-control files used to shoot hundreds of miniature spacecraft elements. Extremely complicated space battle sequences were designed in 3D and analyzed to determine the requirements for shooting each spacecraft. This allowed SPI's motion-control unit to shoot the individual spacecraft with maximum efficiency. Once filming of the models was complete, PLF integrated the component elements of each shot into 3D templates that SPI used to complete the final composites.

PLF's team began by previsualizing each space battle sequence. The previz was then edited into animatics so that they could be viewed and approved by the filmmakers in a cinematic context. Once the animatics were approved, PLF created a 3D "virtual stage" for use in testing various stage configurations and in planning camera movement, lighting, and the position of each model in precise detail. PLF supplied SPI's motion control unit with custom move files for each spacecraft, detailed diagrams of the staging, and reference videos for each shot. This reduced the time required to configure setups and program rigs, and virtually eliminated the need to design moves or try different configurations on stage. With decisions about shot composition and motion out of the way, the film crews were able to focus on maximizing photographic quality.

After the elements were shot, PLF prepared "slap comps" of telecined motion-control elements, creating the first accurate views of the fully assembled shots weeks before they were available at film resolution. In addition, PLF imported the camera data back from the motion-control units and used it to align 3D scenes with the slap comps. PLF delivered the aligned 3D scenes to SPI for use as templates in integrating the photographic and computer-graphic elements.

The previsualization process employed in "Starship Troopers" proved to be an effective replacement for established motion-control production methods, which are costly, inefficient, and ill-suited to the complexity of today's visual effects. Since shots begin and end as single 3D scenes on a computer, motion-control shoots are transformed into a process analogous to rendering in computer animation. In doing so, the previsualization process brings a high degree of certainty to what has been a very uncertain process. That is good news for producers in their desire to prepare accurate budgets and for filmmakers in their pursuit of ever-more-dazzling visual effects.

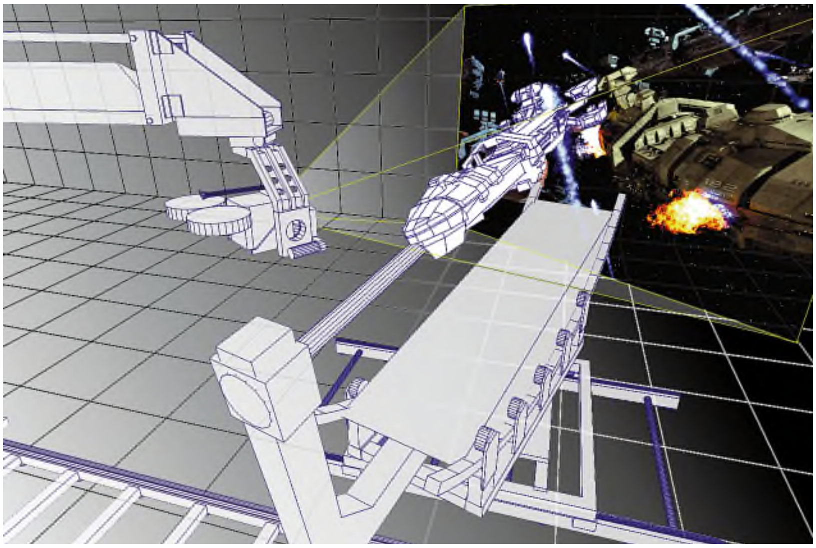

Virtual Fishtank

Henry Kaufman, Brian Knep, Aubrey O. Francois, Tinsley A. Galyean, Stacy Koumbis

Nearlife, Inc.
147 Sherman Street, Suite 102
Cambridge, Massachusetts
02138 USA
henry@nearlife.com
www.nearlife.com
www.tcm.com

Collaborators
Darrin Bascome, George Bird, David Friend, Sheri Galyean, Darrin Jewell, Pat Porter, Thomas Porter, Mitch Resnik, Karen Sideman, Oliver Strimpel, Scott Yu, David Zung.

The Virtual FishTank is a permanent interactive exhibit at the Boston Computer Museum designed to enhance public awareness and understanding of complex decentralized systems and emergent behavior. The exhibit was designed and developed at Nearlife Inc. in collaboration with the MIT Media Lab and the Computer Museum. To provide an engaging and entertaining view of how individuals can follow simple rules that lead to emergent group behavior, the exhibit focuses on the schooling behavior of simulated fish.

Environment

The overall feel is of being submerged at the bottom of a virtual undersea world inhabited by giant versions of kitschy fish tank objects. There are the obligatory diver, treasure chest, clam, and sunken ship. There are five different kinds of fish (two types are predatory), whose styles are inspired from actual fish but are more cartoony and robotic to highlight the fact that they are simulated. The fish swim freely between twelve 1024 x 768 rear-projected screens spanning a region of over 400 square feet with almost 10 million pixels. Recent advancements in processor and graphics accelerator speed allow the entire exhibit to run on 21 home PCs.

Architectural details and interactives that cross between the real world and the fish tank world create a sense of illusion that visitors are actually in an undersea world. A wheel is connected to a pipe that goes into the wall. The pipe continues into the virtual space to an opening where virtual food pellets are released. When a visitor turns the wheel, food is released, and the fish respond if they are hungry. The fish can also react directly to visitors by means of a digital camera that detects motion using simple vision techniques.

The user interface elements were designed to highlight the system of rules that governs the behavior of each fish. At four touch screens, visitors change the parameters of the fish rules, such as the strength of attraction or direction matching, and then watch the schooling patterns that emerge as each fish responds. At three more touch screens, visitors build and test their own fish in a miniature tank and then release it into the big tank. As visitors change the fish's rules, its appearance also changes. For example, making a fish very hungry will give it a large and fierce mouth.

System Architecture

The fish tank software is written on top of Nearlife's Directable Character architecture, which allows real-time, interactive, networked, autonomous character behaviors and extends the work of Blumberg & Galyean[1] and Perlin & Goldberg.[2] It is written entirely in Java, using the latest technology from Sun Microsystems, including Java Development Kit 1.2, Java3D, Java2D, Swing, and the Java Shared Data Toolkit.

In the Directable Character system, every object can be a character with its own behaviors, and each character can be distributed over a network. In this exhibit, the fish and environment characters (e.g., bubbles and food) are created and simulated on one machine, the server. This central machine sends out dynamic state changes to the rendering machines, including low-level degrees of freedom, such as position and velocity, as well as high-level behaviors that do not affect the simulation (e.g., "bite" or "look scared"). The user interface machines create their own local characters (e.g., buttons and sliders), and send state changes to the server. When the server receives these changes, it adjusts the appropriate simulation parameters.

To reduce network traffic, the tank is divided into regions. Each rendering machine only receives messages about fish that are visible. The system automatically handles transitions when a fish crosses between regions.

The simulation time-steps are decoupled between the server and rendering stations, so dead reckoning is used to calculate rendering frames in-between networked updates. A Hermite spline is constructed to animate the fish between where it is now and where the server says the fish should be. If the fish reaches the goal location before getting another update, its position is extrapolated.

The schooling rules are based on Reynolds[3] with extensions for steering forces, predator/prey behaviors, and environmental behaviors such as attraction to food. To make the behaviors appealing and understandable, the motion simulation is augmented with hand-animated sequences. For example, a predator attacking a fish lunges and triggers a bite animation. The prey, sensing an approach, tries to escape and triggers an exaggerated fear reaction. In addition to their entertainment value, these character animations help display the internal state of the fish rules, making it easier for visitors to make the connection between rules and behavior.

References

1. B. Blumberg, T. Galyean. Multi-level direction of autonomous creatures for real-time virtual environments. Computer Graphics (Annual SIGGRAPH 95 Proceedings). 47-54, 1995.
2. K. Perlin, A. Goldberg. Improv: A system for scripting interactive actors in virtual worlds. Computer Graphics (Annual SIGGRAPH 96 Proceedings) 205-216, 1996.
3. C. Reynolds. Flocks, herds and schools: A distributed behavioral model. Computer Graphics (SIGGRAPH 87 Proceedings) 21(4): 25-34, 1987.

Session: Emerging Technologies

Solve et Coagula: mating man and machine

Stahl Stenslie
c/o Telenor R&D
Instituttveien 23
2007 Kjeller, Norway
stahl@rosa.nta.no
televr.fou.telenor.no/stahl

Solve et Coagula is primarily an attempt to give birth to a new life form: half digital, half organic. Through a multisensorial, full-duplex sensory interface, the installation networks the human with an emotional, sensing and artificially intelligent creature; it mates man with a machine turned human and everything that goes with it: ecstatic, monstrous, perverted, craving, seductive, hysterical, violent, beautiful.

Solve et Coagula is a post-human life form. It presents the emergence of a new species, a bio-cybernetic symbiosis, transforming the conception of being human by networking man with a sensing and emotional machine. Through a bio-cybernetic interface, the installation extends the computer's logic and intelligence into a human domain of dark desire.

Through a corporal and sensual symbiosis of an intelligent, interactive three-dimensional creature and human user(s), Solve et Coagula questions the possibility of a post-human and a post-biological life form.

Interaction
The installation interacts with the user through body and voice:

- Body: The bodysuit worn by the participant serves as an intelligent, two-way communication interface to the creature. It provides tactile stimuli so that the creature can touch the participant's body, and built-in bio-sensors through which the creature can sense the body condition of the human.

- Voice: Just as the creature expresses itself vocally to the participant, so does it in turn respond to the sound of the participant's voice. It attempts to analyze the pitch and tone of the participant's utterances as some kind of emotional feedback, and then responds to the participant as it deems appropriate.

Appearance
On entering the installation, the participant steps inside a massive ovoi shell of metal arms. The participant is dressed head-to-toe in the tools needed for interacting with the artificial intelligence: a lightweight bodysuit, a microphone, and a head-mounted display. Video beamers project a view of the creature's constantly changing body on projection surfaces placed around the installation.

The creature speaks to the participant through moving, three-dimensional sounds projected from the eight-channel sound system surrounding the installation. Its "voice" is a bizarre, mutant-like combination of organic noises and distorted samples of the human voice; it changes constantly in response to the participant's input.

The participant wears stereoscopic viewing glasses through which she is visually immersed inside the computer-constructed 3D reality of the creature's body. The creature presents the participant with organic "body parts" as representatives of its various emotional states. The creature brings the participant to each body part through a vein- or intestine-like labyrinth of tunnels.

The tactile body suit worn by the participant provides stimuli of pressure and vibration, so that the creature can touch and manipulate the user's body. The creature can also sense the body condition of the human through bio-sensors built in to the suit.

History
The project was initiated by Knut Mork and Stahl Stenslie. It was realized together with Karl Andersygard and Lars Nilsson. The project questions what happens when the machine turns human and the human turns machinelike. Solve et Coagula attempts to establish broader channels of communication, and even union, between the two. The work addresses itself to anyone concerned with the problematic sides of human-machine interaction, and with the estrangement caused by machines.

Technical Requirements
The installation uses a high-end graphics computer (SGI Onyx) and a custom-built bodysuit controlled by proprietary interface technology. All software was written by the project group.

Diorama

Karrie Karahalios
Massachusetts Institute of Technology
20 Ames Street, e15-466
Cambridge, Massachusetts 02139 USA
kkarahal@media.mit.edu

Diorama is an augmented-reality system that enables one to populate a physical space with three-dimensional objects. It consists of two parts: authoring tools and video tools.

The authoring tools provide a means for placing and manipulating objects in the physical space. The authoring kit also provides a means for attributing simple behaviors to objects. For example, a virtual object might be created as a guide to lead one to a certain location; a virtual tree might be created that grows every time you make a bank deposit.

The viewing system consists of a hand-held lens. This currently exists in the form of a laptop. It is equipped with a small ccd camera and tracking sensors. As users move the lens, they see the physical world with an overlay of three-dimensional objects. These objects are rendered accordingly as users maneuver the lens. The conceptual mapping for Diorama is simple. A virtual camera is directly linked to a physical camera; a virtual coordinate system is directly linked to a physical coordinate system. However, there are technical issues that make this mapping a challenge in practice.

The position of the camera must be known to a high degree of precision in order to define the model of the virtual camera and render the objects in their correct location and orientations. The calibration is currently done using the orientation sensors on the lens. Position is currently handled at certain nodal points. Originally, we began position sensing using an IR sensor grid. We have since moved on to consider an RF approach analogous to GPS.

Within the realm of the current system, we are experimenting with different rendering techniques as the user moves the screen. In one mode, as the screen moves, the objects all appear in wireframe mode. As the user settles into a viewpoint, the objects slowly evolve into complete structures.

We are also experimenting with certain perceptual issues with regard to rendering. This work involves blending an abstract representation of the physical world with the virtual elements. Using such methods, we can impose constraints in an environment with many degrees of freedom.

Some conceptual theories we are exploring with Diorama include issue such as privacy, ownership, value, and communication. Notions of publicity and privacy will have to be rethought to account for this new dynamic space. The boundary of a physical wall does not stop virtual perception. Public and private spaces will overlap at different levels as the complexity of perceptual filters advances.

Ownership of objects and of physical and virtual space creates interesting dilemmas. Perhaps one has write permissions in one's personal office. What rights exist in public hallways? Can anyone who can see these objects move them or remove them completely? The "value" of the real estate of these layered spaces is also novel. Virtual objects might have no correlation whatsoever to the value of the physical space they occupy.

Also, Diorama allows people to use communicative objects to provide visual information as well communicate with each other in this medium.

We are very familiar with the environment we inhabit. We can differentiate indoors from outdoors. We know we cannot walk through walls, and we expect that only one entity can occupy a specific point in space at an instantaneous moment in time. What would happen if one or more of those variables ceased to be true?

The interest in creating a "parallel" world ranges from the ornamental to the behavioral. It is an exercise in placing the fluidity of computer programs in a physical space.

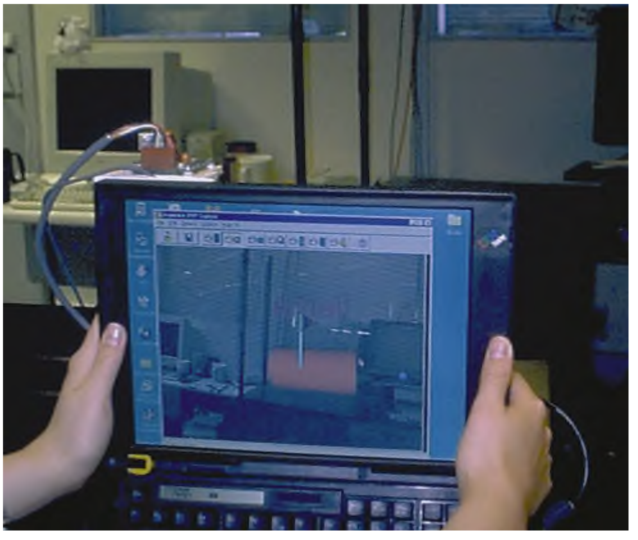

Exquisite Fun: A Digital Sketchbook

Peggy Reinecke
10314 Rockbrook Road
Omaha, Nebraska 68124 USA
pardfa@home.com

Rebecca Hermann
5610 Hickory Street
Omaha, Nebraska 68106 USA
jmbbah@aol.com

How do you use a computer as a sketchbook? It helps to identify how and why artists use sketchbooks. Long before there were computers, most artists kept sketchbooks for a variety of reasons: sometimes as a record of an image or event or thought that impressed the artist; often as a place to practice drawing with both the eye and the hand; frequently as a scrapbook in which to keep photos, postcards, and sketches made on napkins. Sketchbooks become reference manuals in a very real sense: artists, looking for inspiration, often page through their sketchbooks to find an idea, or to refresh the memory of an experience recorded there. From these sketches, other work develops. Even today, the intimacy of the sketchbook prompts art lovers to purchase them in order to view the workings of the creative mind before any censoring or editing takes place.

Three years ago, frustrated by frequent upgrades to the software we use and teach, and unable to meet on a regular basis due to conflicting schedules, we decided to design a project as an antidote to those conditions. The project became a digital sketchbook that exists as a series of files on a shared disk. “Exquisite Fun” borrows its collaborative method and name from the Surrealist technique “Exquisite Corpse," in which one artist began a poem or image, folded the paper to conceal what had been done, and passed it along. The next artist would add content, not knowing what was already drawn or written. According to the Dictionnaire abrege du Surrealisme 1938, the example which has become a classic, and to which the game owes its name, was the first sentence produced by this method: "Le cadavre exquis boira le vin nouveau." (The exquisite corpse will drink the new wine).

In the digital version of “Exquisite Fun," we use the computer and pass a disk back and forth to each other. The rules are pretty simple: when you have the disk, create an image file and a text file, and pass it back. Although we roughly standardized the image size, the choice of image and text content, and the software used, are completely at the discretion of the artist. Often the artist incorporates some portion of the previous image into the current image, but not always. Now almost 100 images and their companion text files exist. This project began as a way to master new software and share creativity with each other. The astonishing co-product is content. A wealth of original images wait patiently on the disk for future use.

As instructors, we noticed that, once hooked they're into their computers, students and professionals rarely interact with each other. This behavior runs counter to traditional artists, who thrive in a community of active exchange. By designing “Exquisite Fun” as a collaborative project, the participating artists challenge and inspire each other. “Exquisite Fun” continues to energize us even after nearly three years of digital sketching.

Creating images in the digital ether poses new challenges to artists trained in traditional media. Although the fluidity of this method allows creativity to flow despite crowded schedules, occasionally visual and tactile artists crave a more solid product. The “solid” form this content takes helps push the concepts of “output." These same images and their companion text files have been a single dyptich, have filled an entire gallery with picture rails and traditionally mounted prints, have danced through an animation, and will be a garden installation this summer. This varied output challenges curators and audiences to question their concept of permanence and the preciousness of the object. With “Exquisite Fun,” the emphasis lies on the creativity these physical forms contain for the moment. A new kind of permanence is afforded through digital technology, a permanence less like a bronze sculpture and more like the permanence of written music with its relationship to performance.

This work exists at a convergence of traditional and digital, of actual and virtual, of relative space and time. When we speak of”Exquisite Fun," we sometimes refer to it as a journal, or a scrapbook, or any other form of documenting change over time. We hope to communicate the excitement this process has ignited about the way the computer makes creativity visible and, to some degree, retraceable.

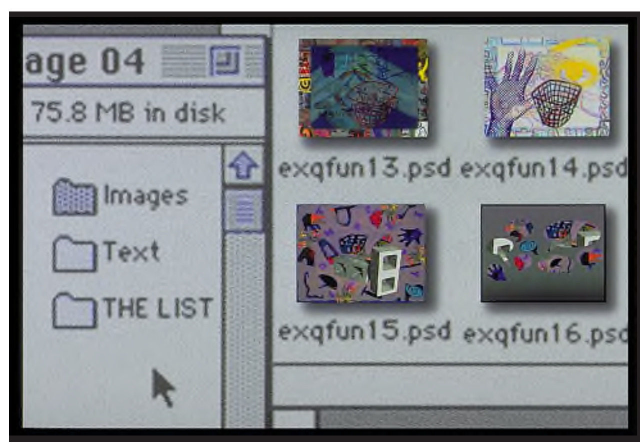

Electronic Remapping: Body Augmentation in the Electronic Age

Bill Hill
Florida Southern College
9506 Waterford Road
Jacksonville, Florida 32217 USA
hill@wvwc.edu

Almost all cultures, past and present, practice some form of body modification. The oldest human remains found to date, the mummified body of a man frozen in the ice of what is now the Italian Alps over 5,000 years ago, had tattoos. In the electronic age, prostheses of all kinds are embedding themselves inside our culture, everything from internal pacemakers to external limbs. As our collective culture and our physical bodies coalesce with technology, there is a essential and traumatic remapping of physical and psychological networks. This remapping is the grounding for two recent electronic installations: Apparatus 3957 and Pre-Programming.

Preprogramming is an interactive installation where the viewers/participants are taught to remap linguistic symbols into letters, words, and meanings. The installation consists of a physical apparatus and a virtual interactive interface. The user actively consumes electronic technology while revealing the fabrication of electronically published information. When cryptic buttons on the apparatus are pushed, a corresponding letter is revealed on the projected interface. The user must decipher the keypad and decode the electronic information. The heart of this work revolves around the ease of electronic publication and its separation with an assumed knowledge, and the precondition of machines to force humans to adapt to their methods.

The role of technology entails not the extension of the body and its images, but more fundamentally, the saturation, replication, alteration, and creation of the organic processes of the body, if not the body itself. Much like amputees, who remap their muscular networks to control their prostheses, users/participants allow the technology to alter the function of their own language.

As the flow of information accelerates, so does the capability of the humanmind to process that information. For nearly 400,000 generations the human species was accustomed to the speed of nature. It understood the complexities and nuances of the natural world. With the Industrial Revolution, the world and many of its parts began to move at mechanical speeds. Interaction with assembly lines began to force humans into an existence controlled by a new speed. Today, with the constant bombardment of electronic symbols, we have developed the process of moving edited images rapidly through a passive brain, which is so different from active information gathering, whether from books or newspapers or walks in nature. As a result, people could become more passive and have less understanding of world events even within an exploding information environment.

With television, for example, the longer one watches the more likely the brain will slip into alpha level: a slow, steady brainwave pattern in which the mind is in its most receptive mode. Because the screen is small, there is little eye movement, so the image comes in whole; the flickering of the screen produces an almost hypnotic trance; and the image comes at and into the viewer at its own speed. In this unconscious state, information can be placed directly into the brain.

The human species is confronted with a perpetual universe that moves much faster than ordinary life. As viewers, we sit passively in front of the "tube" witnessing fractured images that are accelerated and condensed. These types of experiences alter the process by which humans filter information and read the outer and inner world. Not only do we begin to have difficulty staying awake when exposed to the tempo of the natural, but we also develop the inner thought processes of an electronic signal. The memory begins to assemble experiences as if they were viewed in a fast-paced music video. We begin to recognize ourselves and our experiences in the manner depicted before us on television.

Using classical conditioning, Apparatus 3957 acts a biomechanical device to alter the human reaction toward electronic technology. As a kinetic/video installation, Apparatus 3957 consists of a motor-driven device that controls the movement of a human specimen; a single-channel video with four suspended monitors, which documents a previous session; and a body apparatus containing a LCD panel.

In theory, the human specimen will be conditioned by the machine to become physically nauseous when confronted with a computer's operating system. By hooking humans into the apparatus and locking them into a circular motion, forced upon them by the machine, their bodies will produce physical reactions associated with the visual and audio stimulus.

As in our contemporary culture, the human must adapt to fit the parameters of the machine. In this case, the computer assumes the role of programmer of the human specimen.

Interactive Poem

Naoko Tosa
ATR Advanced Telecommunications Research Lab
2-2 Hikaridai, Seika-cho, Soraku gun
619-02 Kyoto, Japan
tosa@mic.atr.co.jp
www.mic.atr.co.jp/~tosa

Interactive Poem is a new type of poetry in which a participant and a computer-generated poet (Muse) create a poem by exchanging short poetic phrases. The basic idea comes from Renga, an old Japanese poetic form that is generated by several people as a combination of short Japanese poems such as Waka or Haiku. This combination has been used since ancient times to express spiritual emotions.

Interaction

The face of the Greek goddess Muse appears on a large screen and speaks short, emotional poetic phrases. Hearing these words, the participant is able to enter the world of the poem and, at the same time speak to Muse. Through this process of exchanging poetic words, the interactive system allows the user and the computer to work together to build the world of an improvised poem filled with inspiration, feeling, and emotion.

Interactive Poem System

The basic form of Interactive Poem is expressed through a simple transition network. To introduce improvisational interaction into the system, we modified the simple transition network and prepared many phrases that can be linked to phrases uttered by the computer to achieve a higher degree of conversation. These phrases are carefully considered and constructed to have phonemic and meaningful relations with the computer phrases. This transition network is put into the database and used to control all processes.

To recognize the meanings of phrases uttered by a user, the system adopts speaker-independent speech recognition technologies based on the Hidden Markov Model algorithm. In addition, the system recognizes emotions. As the basic architecture for the emotional recognition component, the system uses a neural netowork technology.

Generation of Muse's Reaction

The reaction of the computer character to the phrases uttered by the user is expressed through voice and images. For each phrase to be uttered by Muse, multiple utterances with different emotional expressions are stored in the system. By combining the speech-recognition result and the emotion-recognition result, an utterance with specific meaning and emotional expression is selected and spoken by Muse.

In the same way, various kinds of Muse facial images are stored and, based on the recognition results, an appropriate image is selected and displayed. To create natural transitions from one Muse expression to another, 3D morphing animation technique is used. In addition, to express the atmosphere of the world of a poem, the system contains a number of background scenes and displays them depending on the process of dialogue between the participant and Muse.

Virtual Scenery in Broadcast Television: The Time 100 Project at CBS Television

Ann Latham Cudworth
CBS Television
530 West 57th Street Room LE-2=0D
New York, New York 10019 USA
+1.212.975.7612
+1.212.975.7530 fax
ACDesigns@compuserve.com

CBS News is producing a series of television specials in collaboration with Time Magazine's selection of the 100 most important people of the last century. Over the next two years, persons chosen to represent categories such as statesmen, entertainers, icons, industry titans, and scientists will be highlighted with archive materials, interviews and special reports.

In the final show of the series, Time Magazine's Person of the Century will be revealed.

Early in the planning process, the producers decided to use virtual scenery, as it has capacities beyond that of traditional scenery for display of the archive footage, graphics, and video clips. Furthermore, this was a built-in opportunity to develop this cutting-edge technology specifically for use on broadcast television

Design Process

Due to a shortened design development period of only four weeks, we started with several rough prototype sets, built and animated in 3D Studio Max. From this selection, we could move quickly toward the single composite set that supplied all of the necessary parameters for the scope of the show. Over the next two weeks, daily animations were created on the model, as we developed the textures and animated parts that would support the tone and overall visual impact that the scenic required. Decisions regarding animated textures, video feeds, and other real-time performance factors had to be balanced with design characteristics required to accommodate the overall look. When the director joined the team in the third week, further changes were made to the set to accommodate the camera angles and talent blocking. Almost all of the final shots seen in the broadcast footage, were prepared and rendered to a beta test tape and hard copy for viewing by the entire team. This heavy pre-production visualization was considered necessary because we only had one day to complete the shoot.

Installation

After final approval on the geometry and textures of the set, the model was transferred over to the ONYX2 system via .dxf format. The model would ultimately reside in SoftImage and run in the RT-Set system in real time during the taping process. It took several days to link and optimize the hierarchy of the model for its display and control in the RT-Set system. Eventually, the set was in place, running stably in RT-Set at 60MHz.

Our next challenge was in the real world. For two days, we worked with the stage manager and crew to place real platform step units into the chroma key environment and align the virtual set with them. The final result (Dan Rather, walking down virtual steps) brought gasps of disbelief from the control room crew.

Production

With the set working well and the basic talent blocking in place, the focus shifted to the camera moves and what could be developed with the virtual and real cameras. In the studio, we had one camera with a teleprompter, and the RT-Set system was calibrated only with the ability to accept tilt, pan, and zoom moves. So, to "free" the camera, we used the RT-Set system's capacity to hand-off its image from a virtual camera to a real camera. Several shots were combinations of an initial virtual move culminating in a real camera move. Here the teamwork of the RT-Set operator, the director, and the camera operator were of paramount importance.

Continuing Production

During the week of 11-15 May, we went back into the studio for the second segment of this series, Artists and Entertainers, with Mike Wallace as host. As we moved our set up to another studio, we discovered an interesting problem. In trying to replicate the shots we had created in the first segment, we discovered an information gulf between the virtual and real worlds. We needed a way to show the orientation of the virtual world to the real studio space and to accurately locate our real-world objects in virtual space. While we only had a few steps and platforms to deal with in this setting, the time will come when there will be much more intimate connections between real scenery, such as a working door, and virtual scenery (the wall around the door). We will have to be able to load this set again and again with completely accurate each time.

Future Production

Currently, it seems the real strengths of virtual scenery are in graphic display. Showing the Time magazine covers and the montages are the status quo for these systems, so how do we move on to the next level? How do we get more of the real world into the virtual environment

It/I: Theater with an Automatic and Reactive Computer Graphics Character

Claudio Pinhanez and Aaron Bobick
Massachusetts Institute of Technology
20 Ames Street, Room E15-368C
Cambridge, Massachusetts 02139 USA
iti@media.mit.edu
www.media.mit.edu/~iti

The objective of our research on computer theater[2] at the MIT Media Laboratory is to experiment with radically new ways of using computers in theatrical performances. "It/I" is our first publicly performed result: a play in which one of the two characters, It, is not only virtual (constituted by computer graphics objects projected on stage screens), but also controlled automatically by a computer. This reflects our current focus in building automatic and reactive computerized actors able to interact with human actors on sensor-loaded stages and in interactive environments, with no human puppeteer involved in control of the characters during the performances.

Having automatic computer actors on stage enables novel audience interaction. If a character is controlled by computers, members of the audience can re-enact the story of the play in the role of the main characters. In this scenario, a play is expanded from the ritualistic happening of the performance into an interactive, immersive universe to be explored and experienced by the user.

The Play

The computer theater play "It/I" was written in 1997 as a pantomime for a human character, I, and a computer-controlled character,It, in a story about the relations between mankind and technology. The play is composed of four scenes, each a repetition of a basic cycle where I is lured by It, is played with, gets frustrated, quits, and is punished for quitting. In this scenario, It represents the technology surrounding and many times controlling us; that is, in "It/I", the computer plays itself.

It has a non-human "body" composed of CG-objects – abstractions of clocks, cameras, televisions, electrical switches – that are projected on stage screens. It "speaks" through large, moving images and videos projected on the screens, through sound played on stage speakers, and through the stage lights. This contemplates our view that one of the interesting reasons to have computer actors in theaters is the possibility of having expressive non-human characters interacting with human actors.

The Technology

The major technical novelty in "It/I" is the scripting language used to describe the story, the interaction with the human actor, and the behavior of It. The 30-minute interaction between It and I is described in an interval script, a paradigm for interaction scripting based on the concept of time intervals and temporal relationships.[3] Interval scripts were developed to overcome current limitations of scripting languages, including the lack of appropriate ways to represent the duration, complexity, and unpredictability of human action, and to handle delay or failure in the production of output.

In an interval script, each sensing or action event is mapped into a temporal interval, and the designer writes functions describing how to start, stop, and determine the current state of the interval. The relationships among the intervals are declared as temporal constraints. During run-time, the current state of all intervals is compared to the "desired" state (determined by propagating the temporal constraints) and start and stop functions are called to change the state of the intervals that are different from the desired state.

The computer system controlling the It character employs cameras to sense the actions of I on the stage. In the performances of "It/I," we employed a disparity-warped background segmentation methodcharacter[1] that computes a silhouette image of the actor independently of lighting changes or screens in the background.

The CG objects composing It's body were modeled using Open Inventor and rendered in real time in Silicon Graphics workstations. To increase the expressiveness of the CG objects, the system automatically generates sound (by playing MIDI files) whenever an object moves, producing interesting changes in personality and mood. Also, we implemented the mime concept of "clicks." A short increase in speed is added to any movement right after a start and just before a stop.

Production and Performances

"It/I" was produced by Aaron Bobick in the summer/fall of 1997 directed by Claudio Pinhanez, and art directed by Raquel Coelho, with vision systems by John Liu and Chris Bentzel, and Joshua Pritchard in the role of I. The play was performed six times at the MIT Media Laboratory for a total audience of about 500 people. After each performance, the audience was invited to go up on stage and interact with It in a scene from the play.

To our knowledge, "It/I" is the first play ever produced that involves a computer graphics character that is both interactive and automatically controlled. We see the play as an important step in our research on computer theater and also as part of the continuing work of our group toward understanding and developing technology for story-based, interactive, immersive environments.

References
1. Yuri Ivanov, Aaron Bobick, and John Liu. Fast lighting independent back ground subtraction. In IEEE ICCV'98 Workshop on Visual Surveillance. Bombay, January 1998.
2. Claudio S. Pinhanez. Computer theater. In Proc. of the Eighth International Symposium on Electronic Arts (ISEA'97). Chicago, September 1997.
3. Claudio S. Pinhanez and Aaron F. Bobick. "It/I:" A theater play featuring an autonomous computer graphics character. MIT Media Laboratory Perceptual Computing Section Technical Report #455, January 1998.

Designing with Words: a Model for a Design Language in a MOO

Anna Cicognani
University of Sydney
anna@arch.usyd.edu.au

The idea of linguistic performance is intuitive among whoever uses computers more or less extensively: typing commands makes computers do things. One of the hypotheses behind my research is that, as commands make computers perform actions, some commands can perform design actions. In other words, since computer systems and their components (operating systems, programming languages, software, interfaces, and so on) use languages of various kinds to transmit/transform information, this information must be, somehow, the result of a linguistic performance.

I am looking at the constructive and performative aspects of language so that design in text-based virtual worlds (VWs) is achieved. My main claim is that a language for design in text-based VWs can be defined, through an analogy between natural language speech acts and computer commands. The analogy is based on the assumption that language can perform in text-based VWs, since they are, indeed, text-based, and that this linguistic performance is implemented through commands: specific commands for design can be defined to make them perform design-oriented activities.

I have looked at text-based VWs, like MOOs, to identify design parameters and hypothesise a design characterisation for virtual entities. Design in MOOs is of a different nature from real-life design. Starting from a model of structure/function/behaviour of real-life objects, as developed by some design theories, I observe that the structure characterisation becomes irrelevant in a MOO. For physical entities, structure is based on their physical and geometrical properties. Virtual entities do not have to respond to a fixed geometry; they can be organised in a topology of relationships between areas and other entities, and they do not need to obey geometric rules.

The other characteristics of virtual entities can be outlined as:

- Referent, or what an entity refers to (in terms of metaphor) in the physical world, when named in a certain way. A set of expectations on entity performance arises from it.
- Actions, or the permanent (but reversible) changes that entities and commands can perform on the environment.
- Reactions, or the temporary outputs that entities and commands produce.

These three characteristics provide enough information for the design of virtual entities (products). Specific words (processes) can then be chosen in order to perform design in a text-based virtual environment. As speech acts perform actions in real life ("I pronounce you man and wife"), so design speech acts produce design actions in VWs: Speech Act Theory (Austin, Searle) acquires a specific importance in formulating a design language for text-based VWs. A design speech act will be effective if the command syntax used corresponds to an appropriate change of characteristics on the entity considered. I have selected some words for design and some relative scenarios where these words can be used, and I propose some processes by which design is possible.

Developing the Interactive First Person P.O.V.: Using Characters as a Sensory Lens

Ella Tallyn and John F. Meech
Hewlett Packard Laboratories
180 Ashley Down Road
Bristol BS7 9JZ United Kingdom
ella@hplb.hpl.hp.com

The use of characters as a sensory lens is intended to enable creators of story-based, interactive experiences to develop characters in greater depth. Recent story-based computer games such as The Neverhood and Blade Runner demonstrate sophisticated use of changing pace and viewpoint during the game experience. In more traditional media, there are times when the audience is encouraged to soak up subtle information about the atmosphere, surroundings, and characters' personalities and emotions. This technique aims to exploit that area of experience. Creating psychologically complex central characters has not been an important focus for many games and has proved difficult when attempted. A satisfactory relationship between the player and the character they play is difficult to achieve, especially in a first-person-point-of-view situation. Often this is resolved by creating an unassuming "puppet-like" character that does not encroach on the player's feelings or ambitions. In the case of first-person viewpoints, all signs of characteristics are removed. However, this will not suffice for developing rich stories in which complex and individualistic characters are portrayed.

This technique is demonstrated in a scene in which two characters are in a doctor's waiting room. The room contains various objects such as posters on the walls, children's toys, and a skeleton hanging in the corner. One character is feeling agitated, depressed, and full of anxiety, knowing there is going to be bad news. The other character is happy, having just had some good news, and is waiting for a friendly chat. You, the audience, can see out of the eyes of either character at any point during their five-minute wait.

Time elapses, and a few minor events occur. You have the ability to direct each character's gaze in turn, but the characters will only focus on items that interest them. You can only see each character through the eyes of the other. It is hoped that by the time the doctor calls them in, you will have learned something about their personalities and situations. The interaction in this context is seen as another channel of information for conveying mood and mind-state of the characters. The results of your direction will tell you something about the character (e.g., a character is tired and depressed, so the eyes are not raised much above the horizon). Creating viewpoints in this way also provides a rationale for editing sensory information to create intrigue (e.g., one character will not look at certain aspects of the room and becomes fixated on a particular poster. The technique forms a subtle dual control between the audience and the character.

The real test of this piece will be whether players are irritated by or interested in the way their interactions are restricted to the character's behaviour. As with many experiences of this type, much of the response depends on what the audience is expecting to happen. This is meant to be a gently interactive piece, but the technique could be transferred to more frenetic scenes containing many characters. This could aid authors in directing the players' attention to crucial story action at specific moments. Further developments will include exploration into character's body movement and speech.

Personal Computers as Performance Instruments

Lucia Grossberger-Morales
103 Natoma #5
Santa Barbara, California 93101 USA
lllucia@well.com

Personal computers have the potential to be powerful collaboration tools in performance art. They can be used to generate interactive backdrops, or to create virtual characters that can interact with the performer, even if in a limited way.

In 1979, when I got my first personal computer, what fascinated me was the possibility that computers could create the visual equivalent of musical instruments by performing in real time. In 1997, I designed CyberVato Performance Instrument for Guillermo Gomez-Peña and Roberto Sifuentes for "Mexterminator," an interactive performance piece that was presented in San Francisco. The CyberVato Performance Instrument was designed as a visual backdrop, but it also included small animated vignettes that gave insight into the CyberVato character and, additionally, contained animations that Cybervato and the audience could interact with on stage.

My performance instrument designs are meant to be flexible, giving the artist the freedom to make every performance unique. One of the big advantages of computers over traditional video is that it is not linear, so the performer isn't "driven" by the video. "Go Zapata, Go" is designed with this in mind. I created this piece along with ZapNet, a collective of students at the University of Texas. In this performance instrument, there is a library of animations that contain information about the Zapatistas and focus particularly on their goals and their relationship to the Internet. The graphics are used to augment and give emotional impact to what the performer is saying. But the performer is free to use the visuals in whatever way feels most appropriate.

Rachel Rosenthal, the grand dame of Los Angeles performance art, had a pet rat for many years. Using a personal computer, we will virtually resurrect her rat, and she (or someone in her studio) will interact on stage with the virtual rat.

Working as a computer artist in collaborations for the last 20 years, I have sometimes been the most technical and at other times the least technical collaborator. Each collaboration is unique. The most important qualities that I have found for an effective collaboration are: having a shared vision, feeling comfortable to voice your opinion, patience, and the ability to really listen.

The world of performance art is more experimental and therefore more open to using personal computers. The stages are generally smaller or less defined. Whereas in a large traditional theater setting the lack of resolution and speed of a personal computer might prove inadequate, in performance art the personal computer may be an exciting option, not only in the generation of visuals, but also for its interactive possibilities.

Digitrama

Takehiko Nagakura
School of Architecture and Planning
Massachusetts Institute of Technology
takehiko@mit.edu

Digitarama Development Team
Takehiko Nagakura: design, control software, rendering
Haldane Liew: working design, rendering
Hyunjoon Yoo: construction
Sungah Kim: geometric modeling
Shinsuke Baba: preliminary model
Bryant Yeh: video editing

Digitarama has been designed to interactively display architectural form in multiple projections, to operate in response to the physical actions of the viewer, and to expose and visualize its mechanical principles for intuitive comprehension. The apparatus was originally developed for a museum display of a computer-generated reconstruction of the Hagia Sophia, Justinian's great church in Istanbul. The building is modeled in its pre-Islamic form before the earthquake of 557.

Understanding a complex architectural form like the Hagia Sophia often requires analytical representations in various projected views such as plans, sections, interior perspectives, bird-eye views, and abstracted block models. These projections must also be interconnected to permit assembly of a cohesive mental image of the building. With the help of interactive computer graphics technology, this apparatus displays such representational views simultaneously on a projection screen and a flat-panel display, each attached to one end of a rotating arm. A small physical model is also placed at the center of this arm, which rotates vertically and horizontally around the model.

When the viewer rotates the arm, its spatial relationship to the fixed model changes, and appropriate images of the building are retrieved and displayed. That is, the flat-panel display close to the viewer displays an exterior perspective of the building with its view angle adjusted to that of the viewer toward the model. The projection screen at the far end of the arm displays an interior perspective with its view angle adjusted to that of an imaginary camera inside the model toward the direction of the screen.

Greek scenography depicted the life-size cityscape on the backdrop of the stage set, and re-invention of perspective in the Renaissance took place in front of the real baptistery in the Piazza del Duomo, where Brunelleschi demonstrated his perspective drawing of the baptistery. Here, we see comprehensible scale and positional relationships among the viewers, the image, and the viewed objects, which sensationalized the viewing experience more than just projectional accuracy. The inspiration for Digitrama originates from a desire to regain the proximity between the object and the viewer, which tends to get lost in images of computerized presentation systems.

The operating principle of the apparatus is quite straightforward. The rotation angles of the arm are sensed by pen tablets, and computer software computes view angles from them, retrieving appropriate images, and then sends output signals to the flat panel display and LCD projector. To keep the rotating arm light, the projector was placed underneath the arm's supporting axis and a mirror was attached to this axis in order to reflect the projection beam to the screen. All mechanical parts of the apparatus are purposely left exposed so that viewers can intuitively understand the process of generating output images from user-initiated rotational interactions.

Hagia Sophia is a building with a complex exterior form and a dramatic single interior space. This custom-made apparatus is capable of displaying an interior perspective from a fixed viewpoint and an exterior perspective around a fixed target. It is well suited for interactive graphic representation of architectural form in an exhibition setting.

Las Meninas: The Articulation of Vision

Hisham Bizri
Electronic Visualization Lab (M/C 154)
University of Illinois at Chicago
851 South Morgan Streeet, Room 1120 SEO
Chicago, Illinois 60607-7053 USA
+1.312.996.3002
+1.312.413.7585 fax
bizri@evl.uic.edu

Created by:
Hisham Bizri, Andrew Johnson, and Christina Vasilakis

Las Meninas is a virtual reality artwork created for the CAVE. The painting of the same title, "Las Meninas" or "The Maids of Honor" (1656), by the great Spanish painter Diego Velazquez, challenges the viewer with its allegorical subject matter and enigmatic mise-en-scene. From the outset, the viewer confronts the artist's canvas, which is forever hidden from view. The viewer desires to see what is hidden and at the same time witnesses a scene that carries within itself multiple allegorical meanings: the mirror in the black frame that reflects the half-length figures of King Philip IV and Queen Mariana but nothing else in the room; the magical stillness of the room and the people in it, as if photographed, forcing a belief that the viewer is actively present at the scene; the painter himself whose "dark form and lit-up face represent the visible and the invisible" (Michel Foucault); the imaginary space outside the picture frame where the painter and several other characters are looking, each from a different point, at the sovereigns, who are in theory standing next to the viewer; and so forth.

Thus, the allegorical subject matter and enigmatic mise-en-scene work together in Velazquez's painting to dramatize the "inner focal point" of the realm of the painting and the "outer focal point" of the realm of reality. The viewer is at once "seeing" and "being seen." The vision, therefore, is no longer fixed on a vanishing point, but is now dispersed over multiple planes of form, function, and subjective meaning. The painting raises questions about the nature of representation and subjectivity in a unique way rarely ever matched in the history of visual art.

In the CAVE, when "Las Meninas" the painting becomes Las Meninas the virtual reality, the viewer is able not only to solve certain problems pertaining to the nature of representation and subjectivity, but also reflect on further questions. It is important to point out here that the painting's original size (10 feet square) is the same as that of the CAVE. The very theoretical questions the painting raises become tangible and empirical once placed within the boundaries of VR. In other words, the painting's fixed and traditional problematic of representation and subjectivity takes on a dynamic and physical aspect once the center of vision is dispersed in the medium of VR.

Illusion in art is a complex topic, and each era has its own limitations and paradigms for rendering reality. As inventors of VR artworks, it is important to invent a language that defines the way our new tools of production operate and shape the future of art. In Las Meninas, illusion in art manifests itself as we react, feel, and think in front of the cryptograms of its virtual worlds. In other words, the audience not only witnesses the faithful and convincing
representation of a visual experience through convincing representation, but also the faithful construction and orchestration of a relational model in which the interplay of image and sound trigger in the audience a stimulation to bring about a "second reality." This second reality originates in the audience's conscious and unconscious reaction to the virtual world and not in the world itself.

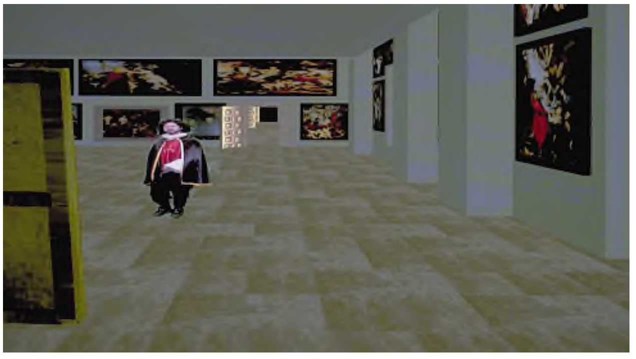

Session: Peep Show and Tell

Gestalt Inhibition Sequence: Digital Hypnosis

Andy McIntire
1449 South Milledge Avenue
Athens, Georgia 30605 USA
amcintire@mindspring.com
www.mindspring.com/~amcintire/gest.html

This project was partly inspired by a need expressed by researchers at Georgia Tech for overcoming a subject's "gestalt" (willing suspension of disbelief) upon immersion into a virtual environment. Subjects being treated for acrophobia (fear of heights) have been facing their fears through VR simulations, but they are distracted by lack of realism. It seems plausible to use VR along with hypnotic suggestion to immerse the subject into a situation, so that the VR serves as a medium of communication and navigation between the conductor and the subject, thereby blending the internal experience with the simulation, rather than relying on technology to muscle out a complete reality. The visualization I have created, which I have dubbed "Gestalt Inhibition Sequence" to thwart academic skepticism, represents a proposed method of psychological immersion into virtual environments using audio-visual wave sequencing linked to biofeedback.

By combining the principles of Eriksonian psychology, biofeedback technology, digital conversion of algorithms into images and sound, and current VR technology, the promise of creating a consistent method of psychological immersion may be realized. Of primary consideration when pushing such a technology as VR is recognizing virtual environments not only as ultimate realities, but also as formation of a nonverbal language to serve as a more powerful avenue to accessing the infinite landscapes of the mind. Perhaps ultra realism in VR has a place similar to descriptive verbal communication (rhetoric). Taking advantage of VR as a means of communication and connection of information within one's own mind or among several individuals could happen sooner provided we have a reliable, quantifiable means of validating perception. Individuals communicating through a VR medium, having been psychologically immersed, could enjoy the efficiency of nonverbal communication, perhaps expressing thoughts on the fly through creative media or interactive behavior. The potential for creating universal abstract communication within a virtual environment could be more readily fulfilled, resulting in new levels of creative progress, provided the perception of individuals could be ascertained. The growth of digital technology is an obvious reflection of our need to overcome our limitations in accessing and connecting the knowledge we already possess. As a symbiotic relationship, we try to bestow life to machines. Therefore, the obvious next step in human-to-machine, human-to-human, or human-to-self interface is to establish a way of verifying what perceptual directory we are operating from.

Description of the Submitted Work

The animation sequence I have created, though recognized as a visualization of the proposed "Gestault inhibition" concept, was created in a more manual fashion, rather than the suggested by-product of a biofeedback-driven algorithm. The animation is the culmination of many sources of inspiration and themes through which I am learning and growing personally. Some of the major themes include: waves; fractal behavior; connectivity; mutation and generational difference; random, fluid, and non-linear movement; scale; resolution and recognition of emerging patterns; layers and transparency both of time and space; and vorticular time, or, focus, and the momentum and gathering energy. Most importantly, I was able to explore composition for proper transition.

The movement is comprised of many generations of submovements layered together, growing slowly, creating a tissue of flurried movement at first, gradually becoming more evidently in sync, and finally united as one movement in the final frame of the sequence. The gradual synchronicity is a constant rate that acclerates the movement overall and contributes to the perception of building energy, while the submovements build to a peak and then gradually subside, providing a less pronounced transition to counterbalance the accelerating synchronization.

The primary tools involved in creation of the sequence were Adobe Photoshop and Adobe After Effects on the Macintosh platform and Alias | Wavefront Power Animator on a Silicon Graphics workstation. Of particular technical interest to me was the use of solid projection for placing each frame of one animation onto animated geometry, which was performed manually one frame at a time for the entire length of the animation. The process of its creation, as well as QuickTime movies, may be viewed online: www.mindspring.com/~amcintire/gest.html

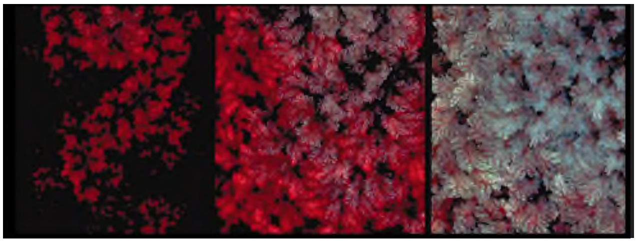

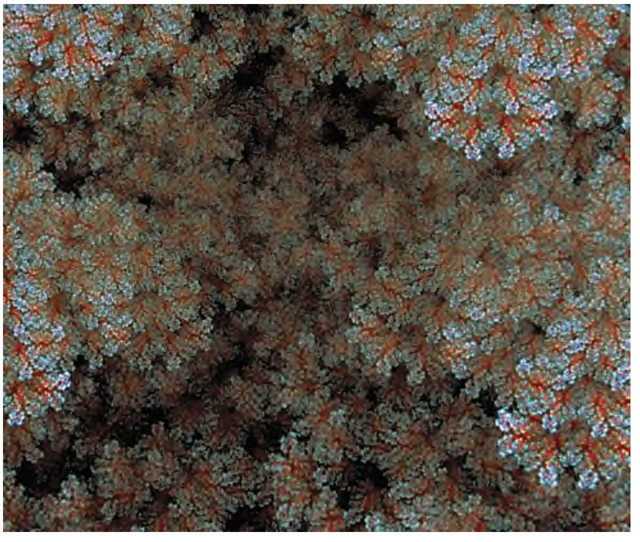

Titanic and Digital Character Animation

Daniel Loeb and Andre Bustanoby
Digital Domain
300 Rose Avenue
Venice, California 90291 USA
bhoffman@d2.com

For the feature film "Titanic," our job was to populate the deck of the ill-fated liner with digital passengers. The animation challenges were formidable, and we approached the problem from many different directions. There was ultimately no single solution. Instead, we used a composite of techniques to arrive at animation that helped tell the story.

From the start, we knew that we were going to have to replicate human beings in 3D and strive for maximum photo-realism in look and performance. Initially, we looked into using performance capture techniques and technology to "digitize" actors' performances. We also figured we could infer extremities by "capturing" behavior of the long dresses and coats that were the style of the day.

Marker counts topped 60 for most of the captures. There were typically two or more performers per capture. This enabled us to preserve a sense of naturalness in the performances. Sample rates were between 60Hz and 120Hz depending on the frenetic nature of any particular performance.

Stunt people were also captured for the sinking sequence. Even though the resultant performances were digital, the performers could still "author" themselves in a dynamic yet controlled way. Two-hundred-foot falls, traditionally very dangerous, could be completed on our stage at only 20 feet and extended digitally. We dressed the "set" with railing extensions, ropes, and deck capstans to give some context to the space and then surrounded the area with cushions and safety bags. The performers could then get some sense of the environment they would eventually be digitized into and be safe doing it.

Although excellent results were achieved with our performance-capture-driven animation solution, we found that the process was too time-consuming and inflexible to satisfy all our needs. Once the captures were processed, skeletons of kinematic chains were generated. This was read into the animation system, and enveloping commenced. Unfortunately, changing this data, from a directorial point of view, was problematic. Unraveling forward kinematics, at that time, was very difficult. So, the decision was made to augment our library of animations via other approaches.

As a first step, we had artists create animations of characters involved in fairly simple activities (standing, talking, leaning at the rail) via a purely key-framed approach. Yet we knew we still had an enormous number of very vigorous animations to create for the sinking sequence and elsewhere, and we felt that they would not easily lend themselves to key framing without reference to some sort of recorded human motion. We came upon the idea of a hybrid system, which we dubbed "rotocapture," that allowed us to achieve the photo-realism of a performance-captured animation with more speed and ease.

This system basically involved: capturing the performances of actors and stunt people with a greatly reduced marker set on a motion capture stage, processing those captures to create animated skeleton-like figures, and bringing those figures into Softimage, the animation package we had selected for the show. We built a very light and efficient proxy model, which we called the "Poopett" (for poop-deck puppet), that we could key frame against the skeletal animation, from which we could derive poses and timing.

A large number of animations were created using this system. Straight-ahead key-frame animation, without reference to recorded movement, was also ultimately used throughout the show, including during the chaotic sinking sequence.

Wiring Cracker: The Mechanics of a Non-Anthropomorphic, Real-Time, Performance Animation Puppet

Mike Morasky
Protozoa, Inc.
2727 Mariposa, Studio 100
San Francisco,California 94110 USA
mike@protozoa.com

Digital Geppetto

At Protozoa we've coined the term "wiring" to describe the art of procedurally connecting data from various input devices to a real-time character or CG puppet. Mapping a human's motions onto a non-humanoid character is one of the more fascinating aspects of wiring. One such character, "Cracker," a genetically mutated "Lizard Cow," was designed and modeled by Bay Raitt and wired by myself for an in-house short animation called "Meat." It's the non-anthropomorphic mapping between the performer and the CG character as well as the techniques used to accomplish this that are described in this sketch.

Beauty in Simplicity

The thing I find most interesting about this character is its apparent simplicity. With only 5 EM sensors producing global position and orientation data for the hands, feet and pelvis, the performer is capable of eliciting a full range of motion and expression from such a non-humanoid character.

To define the position and orientation of an object in space, we generally use xyz coordinates for position and quaternions for orientation. I, however, sometimes like to think of an object's 3D definition in terms of its matrix and component vectors. Viewed in this fashion, a mere five sensors transform into a huge quantity of data that can be utilized using many well-known mathematics techniques such as linear algebra and trigonometry.

Feet on the Floor

The first problem I had to solve on this character was how to get her to walk. To keep her feet from sliding, I stuck to the rule of "feet on the floor," not altering the foot sensor data except for global position scaling. Starting with that rule, you could say I built Cracker from the feet up.

Considering that her knees bend in a backwards, chicken-like fashion, I decided to ignore the knee sensor data and create her knee positions by using a simple inverse kinematics operator.

Because I wanted Cracker to be able to slink low to the ground, taking long strides, as well as walk more upright, I needed to extrapolate two hip positions from the pelvis sensor such that performing her would feel right as well as meet some parameters required by the IK. The solution was to write a "triangle to non-planar quadrilateral" operator. Taking the feet and pelvis sensor data, the operator returns a quadrilateral such that two of the points are untouched foot data, and the other two, the hips, are translated off of the pelvis by parameter-specified vectors. The hips always remain a specified distance from each other (keeping the hip joints solid), and neither ever ventures outside of a sphere originating at its respective foot with a radius equal to the sum of the upper and lower leg lengths. This is accomplished (minus the bells and whistles) by calculating successive frames based on:

1. A spherical interpolation between the pelvis sensor's x-axis and the normalized vector between the feet (percentage of "swing").
2. Orthogonalization of the resulting x-axis with an axis derived by normalizing the subtraction of the averaged foot position from the pelvis sensor's position.
3. A translation defined by the farthest intersection of two rays starting at each hip, pointing along the y axis of the current frame with each hip's respective target sphere.

This allows the performer to be scaled up as large as necessary to allow for long strides without breaking the IK chain, keeping the feet on the floor as well as keeping the character low to the ground. Other adjustable variables include: differing leg lengths, placement of the hips anywhere within the allowable sphere of motion, and non-symmetrical, individual hip placement.

Driving Cracker

Being a mutated lizard, Cracker had to have a very flexible spine and large range-of-head movement. I quickly realized that trying to directly control her head with the performer's head sensor data would be difficult and painful to perform. I decided that how I really wanted to control her head and torso was with my hands, as if I were driving her around.

First, I created a deforming spline based on a hierarchy of nodes with the head perched on the end of the chain. The second-to-last knot in the chain is a child of the head keeping the neck pointed into the base of the skull. To allow Cracker to sneakily dip close to the ground as well as stand up straight in defiance like a cobra, I created a circle in the XY plane through which Cracker's torso would rotate. The amounts of rotation about each of the two axes in the circle are determined by the orientation of the left-hand sensor relative to the pelvis sensor. This is simply calculated by taking the dot product of the perpendicular component axes (z.x, z.y) of the left hand and the pelvis sensors and factoring the result by the desired degrees of rotation per link on the chain for each quadrant of the circle. This gives the left hand a nice "rounded" feeling of control over the torso as well as the ability to adjust the "pose" of the torso depending on her posture.

The head's orientation is then a direct mapping of the right hand's global orientation data times the inverse of the global orientation of the parent of the head in the hierarchy. The eyes, nose, and mouth are wired for performance to be controlled both procedurally and with devices such as a joystick and sliders.

Wag the Lizard Cow

Finally, the tail is another deforming spline whose knots are controlled via a decaying, inverse delay of the pelvis's velocity vector. This creates a nice sine wave effect that makes the tail appear as if it's moving as a whip-like integral part of the character.

Post Performance

Given a good performance, very little refinement is generally necessary. If it is, however, all of the described elements can then be refined by additional expressions, device driven adjustments, channel editing, or key framing.

Starship Troopers

Craig Hayes
Tippett Studio
2741 10th Street
Berkeley, California 94710 USA
craig@tippett.com

Tippett Studio relied exclusively on computer-graphic technologies to model, animate, light, and composite 225 digital bug shots, some with thousands of bugs. This was the studio's first large-scale foray into all-CG production.

Our involvement came early in the development process, enabling us to design the creatures with the specific needs of the various sequences and actions in mind. By working closely with the director, we were able to focus the design of each of the bugs on the role it was to play in the eventual bug-army (infantry, artillery, bomber, and commander).

To pre-visualize the key sequences, animatics (moving storyboards) were executed in low-res form emphasizing the kinetic feel and pacing of the action. These would later serve as blueprints for finished sequences, as well as a concrete, visual reference for the actors and crew working on location.

Because the story called for huge swarms of bugs – sometimes thousands of creatures in the frame at once – no type of animation was ignored. An all-hands-on-deck methodology was adopted, with traditional stop-motion, motion-capture/puppetry, "traditional" key-frame CG animation, and procedural animation techniques all employed to finish the job.

First implemented for "Jurassic Park" as a way for traditional stop-motion animators to apply their talents to CG dinosaurs, the DID (Digital Input Device) was further refined for "Starship Troopers." We built two types of custom-input armatures to match the skeleton of the CG bug. One functioned as a stop-motion puppet and could be moved frame-by-frame and recorded directly into Softimage. The other moved more like a real-time, rod-type puppet, and recorded live performances into Softimage. These devices allowed us to work with the most talented animators even if they were untrained in CG animation tools. The puppeteering technique was especially effective for quickly making many background "extras" for shots involving large numbers of bugs. This input animation, either stop-mo or real-time, was then joined with other Softimage animation, refined if necessary, and then rendered.

In situations where a great deal of control over an individual bug's timing and position was needed, key-frame-style mouse and keyboard animation was the preferred method Animators could do rough-blocking with "floating" hood-ornament style models to establish the basic paths and timing of small-to-medium-size groups of bugs and then go in and refine individual bug movements as required by the shot.

For shots that called for hundreds or thousands of bugs to swarm over the landscape, a procedural approach was used. We used Alias | Wavefront's Dynamation to create particle systems that would swarm toward a particular goal, play follow the leader, or just follow flow fields set up by the animator. At this point, each particle was treated as a sphere to ensure that no collisions occurred, but that the scene would still be computationally tractable. After the overall motion of the swarm was determined, animation was applied. Pre-animated cycles (run, walk, milling around, etc.) were instantiated for each point in the swarm, depending on the point's speed. We further manipulated the particle system so that individual bugs could be marked for special behavior such as falling down or getting shot on a particular key frame.

This procedural animation system also aided our compositing efforts by outputting world or screen-space locations for footfalls and dust elements to blend in with the background plate.

We wanted to take advantage of our animation talent and not let technology guide our work. The sheer volume of animation required for this film compelled us to employ every conceivable animation discipline, from the most traditional to the most technologically current. In the end, we feel we were successful in taking the best each style of animation had to offer and adapting the technology to integrate those various disciplines into the computer-graphic world, where we lit, composited, and rendered the final performances.

Starship Troopers

Walt Hyneman
Sony Pictures Imageworks
+1.310.840.8000
walt@spimageworks.com

Session: Bugs 'n Space

Director Paul Verhoeven presented a formidable task when he described the fleet battles and space settings for Starship Troopers. His vision was one we are all familiar with: something he had never seen before. He wanted a stage to balance the bug wars on the ground, one that would bring epic scope to the film.

Our creative focus and goals were based on execution of the major battle sequences and establishment of the Starship fleet's presence as air cover and troop transport. This involved creation of 122 visual effects shots and utilized thousands of elements and a vast combination of techniques including digital composites, scale models, pyrotechnics, stop-motion animation, practical elements and in-camera effects enhanced, expanded, and supported by digital animation, compositing, and model work.

Our creative inspiration was drawn from historic naval and air attacks but applied a modern point of view and editorial style most similar to car chases. Action was supported by rich and detailed planets, large models, and dramatic lighting. Scenes were designed to give beautiful vistas while keeping the audience pinned and framed within the action. We called this "epic claustrophobia."

The digital needs of the show were numerous and challenging. The crew created the planets, suns, stars, and effects animation required by each scene. The work also included animation of small ships and quick-moving details. Integrating the digital effects work with the live-action and stage elements required an enormous amount of data sharing, match modeling, and match moving.

Planet development led our attack. As our action developed, the planets provided a frame, a claustrophobic limit to our views, and provided the contact with the action below. Motivated by NASA photographs and reference films, we provided each planet with animated atmosphere, storms, and atmospheric shadows, and physical surface detailing such as canyons and mountains. Each planet was rendered in Renderman and digitally composited into the scene.

Engine technology was also a design element, and each flying ship was given a distinctive thruster effect. The Rodger Young class of ships have a bluish comb-like signature, drop ships have blue flame signatures, retrieval boats have whiter flames, other ships have more yellow emissions. In the atmosphere, thrusters also create heat distortion and interactive smoke. All of these elements were data matched, hand tracked, and dynamically fit into the environment, including mattes and interactive effects.

"Bug plasma" created the primary focus of our effects animation team. This concept, developed at Tippett Studio during early testing for the film, was creatively and technically adjusted and adapted by the Imageworks crew for its "space plasma." Bug plasma was the motivating force to everything in our universe. It would cause lighting effects, some of which were built into first-unit photography and had to be replicated, or accounted for otherwise. As such a strong lighting element, foreground plasma had been roughed in by pre-vis and accounted for on stage. Precise action and all background plasma was designed in digital. The bug plasma was generated in a combination of in-house code, Dynamation, and RenderMan.

Our digital animators and technical directors also created the stars, numerous smoke and debris elements, background Rodger Young class ships, and the TAC fighters and drop ships seen maneuvering through the fleet. Procedural code implemented by our technical directors was utilized to assist in the "flocking" and combat maneuvers seen in the distant and fast-moving ships.

Visual Effects by Sony Pictures Imageworks, “Starship Troopers” Courtesy Tristar Pictures and Buena Vista International.

The PDI Crowd System for ANTZ

Luca Prasso, Juan Buhler, Jonathan Gibbs
Pacific Data Images
3101 Park Boulevard
Palo Alto, California 94306 USA
prasso@pdi.com

In 1996, PDI and DreamWorks SKG entered into a co-production agreement to create "ANTZ," a fully computer-animated feature film. Among the many technical challenges we faced were crowd scenes. Since the film takes place in an ant world, many shots involve hundreds of ants and other insects.

A crowd scene is any shot with too many characters to key frame by hand. The amount ranges from a handful of ants standing behind the main characters to over 50,000 ants populating an underground city. So animators must be able to control the motion of background characters with the same flexibility they have on the main characters, yet be able to manipulate them by the thousands when required. To achieve this, we developed two crowd-animation systems: The Crowd Simulator and an extended character setup called The Blending System These systems are written in a powerful procedural scripting language that lies at the core of the PDI animation pipeline.

Since the best generators of character motion are the character animators, the basis of both crowd systems is the motion cycle. Creating cycles is difficult because it requires quality character animation to convey the mood and feeling of the character, yet the cycle must be general enough to work for many different individuals with only slight (or no) changes.

The Blending System adds to traditional key framing techniques the ability to script motion from a library of cycles and create seamless transitions between cycles. It also provides several high-level behavioral controls (such as balance and locking feet) which help adapt the motion to new situations and provide individuality to the character. When assigning motion to differently shaped bodies, the Blending System attempts to preserve the character of the original motion. The animation computed by the Blending System is in a function curve format, allowing the animator to refine the motion with the same level of detail as the completely hand-animated characters. The animator can thus add subtleties that are not possible with an automated system.

The Crowd Simulator's task is to produce realistic motion for thousands of characters. The Crowd Simulator resembles a physically based particle system, except that motion cycles play an important role in how the characters can move within their environment. The simulator drives the choice of a cycle for each character and its overall motion using a combination of physical forces (obstacles, goals, flow fields, intercharacter collision avoidance) and procedural rules (finite-state machines, flocking behaviors). Certain cycles also place constraints on the overall movement of the characters. A plug-in architecture allows the animator to change most aspects of the simulation using shot-specific information.

We use traditional rendering techniques for the crowd sequences along with proprietary techniques we developed to optimize rendering time, I/O, and disk space, which are all critical when rendering crowd scenes.

We have used these systems successfully in every crowd shot produced so far. They have integrated well within the traditional animation pipeline and allowed us not only to produce very complex crowd shots, but to do them throughout the film.

Pontiac Coyote

Indira Guerrieri and Izzy Acar
Industrial Light & Magic

The primary concern of Pontiac and Warner Bros. on Pontiac "Coyote" was to get a level of speed and flexibility in the camera that wouldn't have been possible using traditional multi-plane painted backgrounds, while still retaining the style of a traditional cartoon. Our instructions were to follow the look of the original Chuck Jones layouts as seen in a 1957 cartoon called "Zoom and Bored."

From the start, much concern was raised that CG not "look" like CG. The "Zoom and Bored" layouts had a very graphic style: flat colors, shadows cast by not necessarily consistent sources, off-register lines at the edge of objects, and dark graphic lines defining crags and edges. Our pursuit of this style meant avoiding the traditional "realism" of CG in favor of a flatter, more graphical rendering alternative. We made choices that involved a series of simplifications of the usual CG model.

Softimage was used to model, animate, and render, a departure from the traditional ILM pipeline. We had little time and a fairly simple lighting setup (see below), so that our main worry was to keep the pipeline as basic as possible. We used our in-house compositing software for rough comps that we could look at during dailies.

We split our background renders into many layers (for example, layers for cast shadows, immediate foregrounds, and car reflections). Each layer had a Softimage scene associated with it. All this so as to give greater flexibility to the director and the clients during the final Flame compositing.

On the modeling side, we wanted to create "libraries" of rocks and mesas to be shared among the shots. The hero rocks, as well as plants and props, were modeled as freestanding objects with all-round detail. We modeled patches rather than polygons to give the rocks more rounded edges. For the trademark Chuck Jones lines, we made liberal use of Softimage's Create Modeling Relations. Circles were extruded along splines generated from the surfaces of the models. Scaling and size of the lines were tweaked by modifying the extrusion circles. Should clients change their minds about size, color, or placement of the lines, it would be fairly simple to scale and move these directly on a model rather than have to repaint large numbers of texture maps.

On the rendering side, since we had little concern that the shadows had to be cast from a consistent direction, we cut our patches into areas to which were assigned constant-shaded materials that were either dark or light. No lights were used. We only used CG lights in rendering the cast shadow passes of selected objects. We used no particular shaders for rendering (we looked at Softimage Toonz but decided against it) because most of the real "rendering" work had been done in the modeling stage.

Surprisingly, the issues raised by incorporating a live car into a cartoon environment and making it look like it belonged, were less than we expected. CG backgrounds were match moved to the live plates, as we would have done on any other show. For each car shot, we rendered a separate pass of background reflections on a viewpoint model of the car. These reflections would be screened over the real car in the composite phase. A simplified contact shadow was modeled that followed the car and was itself rendered as a separate pass. To approximate the "realism" of the car we added post-motion blur to the backgrounds with Softimage's stand alone mb. Some of the shots required us to use the viewpoint model as the hero, though the car was so small and fast that we didn't have to worry about achieving a close match.

This sketch covers how we created the number of elements that were assembled on Flame for the final integration of 3D landscapes, traditional characters, and live-action footage. What were the techniques that made the integration successful and what could have been the pitfalls? We show slides of the Warner Brothers layouts, examples of models and rendering, as well as videos of the initial cut, the original Chuck Jones cartoon, and the final commercial.

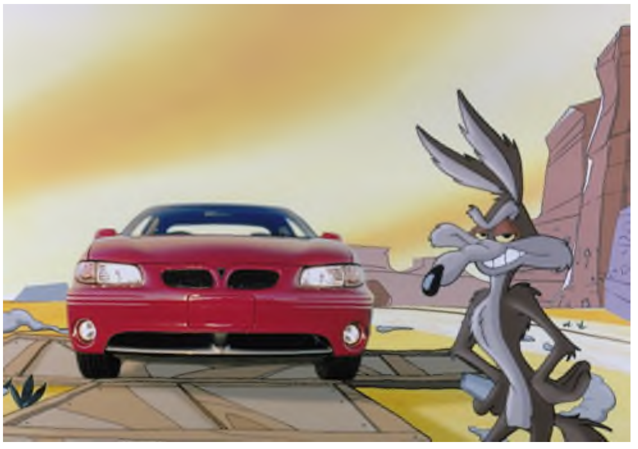

Coca-Cola Factory

George Evelyn
Colossal Pictures

Chuck Gammage
Chuck Gammage Animation

Jerry van de Beek
Little Fluffy Clouds
Pier 29 Annex, San Francisco California
94111 USA
betsy@littlefluffyclouds.co

At Little Fluffy Clouds, our task for this 30-second international commercial, was to recreate a "Rubber Hose" style cartoon animation, a style originally made popular in the 1920s and1930s. This is a look much favored by Colossal Pictures' Director, George Evelyn, who conceived of the spot and drew the original storyboards and model studies. Coca-Cola Factory seeks to transform that early Rubber Hose style into a fast-paced, contemporary Coca Cola commercial. We achieved this by utilizing traditional 2D cel animation harmoniously combined with the latest in innovative 3D digital animation techniques.

The animation and motion for the piece was strongly based on the rhythm of the music from the era. Asyncopated jazz age track, with effects to match, accompanies the animation. Almost the entire piece is animated on doubles to accomplish less of a CGI look and synchronize it more with its original Rubber Hose roots.

The most important issue was the integration of the two techniques of cel and digital animation. Each had to support the integrity of, and not conflict with, the other.We did not want to end up with a modern-versus-traditional, dual-layered look. The production teams, Little Fluffy Clouds of San Francisco and Chuck Gammage Animation of Toronto, succeeded in their goal of integration and perfect registration without a hitch. Successful blending of these techniques results in an interesting image and animation quality where environments are richer and more complex than those of the original cartoons of the 1920s while maintaining the charm and essence of the original style. For instance, real shadows give added perspective. Special code, DIANA, written at Little Fluffy Clouds, by Jerry van de Beek, produced a digital reference file that would print each of the CGI scenes frame by frame, complete with 3D characters and backgrounds. Each scene was then sent to the cel animators as reference for the 2D animation. The use of the digital reference files enables the 2D characters to fully interact with the 3D characters and backgrounds in perfect registration.

The spot was made in black and white. The layers were combined in post production with Flame, and the entire piece was then colorized to give it its final unique look: a 1920s black and white animation with colorization to make certain distinctions in the storyline. The story and thus the feel of the piece is the juxtaposition of two factories and their workforces, one happy and one sad. To quote Director, George Evelyn, In every scene, we gave all the elements a face and a distinct personality – from machines to swinging lamp shades. To create that special look, we made simple but strong shapes in 3D, like the square factories or the wagon, using lots of contrast to create drama. The edges of these models were beveled, not rounded, to give them just the right graphical depiction so that we could reference the shapes in the original Rubber Hose-style. Shadows were very dark and sharp, not soft, to give extra contrast to the different planes.

Extreme light sources were used for the same effect. We used reflection maps as opposed to ray-traced reflections, so instead of reflecting the real environment we could reflect what we chose to onto the texture map. These maps were just simple black and white patterns to create simple graphical reflections.

Once each scene was created and animated, we printed out each frame as a specific reference, using DIANA software. This software ensures that the printed file registers exactly between the digital 3D animation created in Softimage, digital ink and paint in Toonz, and the cel animation. Cel animators then use these frame-by-frame reference files as a guide on which to draw in their characters and animation. This enables the cel characters to interact specifically and accurately with the 3D characters and world. At this stage, interaction between the 3D and 2D animation becomes much more flexible without requiring lots of layering and matting in post. When the cel animators finished their scenes, the frames were sent to digital ink and paint. These layers were then composited together. The 2D and 3D were then combined in Flame.

The spot was rendered in black/white (gray scale). It was then sent to be colorized. Taking the rendering process one step further, different layers and/or different models were rendered seperately and image-processing techniques were applied to them in the final 3D comping stage. In this way, you can add contrast, threshold certain color ranges, darken areas, or diffuse layers to obtain just the effect you desire. Many times, we render the objects seperately from their shadows, or even shading, and comp them together at a later stage to give us more control over the amount of shading, transparency of the shadows, and even the color of the shadow.

A special edition of frames from this popular spot has been selected by The Coca-Cola Company to be sold in animation art galleries.

CPU

Wayne Gilbert
Lucas Digital
48 Magnolia Ave
San Anselmo, California 94960
wayne@lucasdigital.com

The computer is one of the tools of the future and is best used when applying skills of the past. It provides a better way to do cave paintings. It provides a new and arguably better way to communicate ideas to an audience. It's possible that those best able to communicate ideas are those with clear ideas to communicate, by whatever means. It is the idea that is important.

At the heart of our film "CPU" is a comment, not in favor of or against the computer or pencil but in favor of planning. What has come to light during my short time learning how to properly use the computer for animating is that there is no change in the priorities that have held true for decades, not just in animation but in all communication efforts: know what you are going to say before you begin.

With the computer's phenomenal capabilities evolving daily, it is easy at times to get carried away with how something looks and lose sight of what is most important: what is it that the image or images are communicating?

The computer in all of its deserved glory adds an extra element of difficulty to the life of an animator. It becomes a hypnotist that draws you into its beautiful world and makes you believe that your job as an animator is easier with its help. It certainly has made creation of visually stunning images easier, and that might be what is misleading.

Think about this leading question: what if you start to do something before you know what it is that you are going to do?

Back to the computer. It does not think. It will deal with whatever information you give it and actually requires more information than a good animation assistant needs. If you give the computer insufficient information, it will give you the same in return. It will fight you with its bells, whistles, and crashes even more often when you are remiss with supplying information. At times, it will crash with no apparent reason, and yet it calls to us to play with it, before we are ready to play, before we have cleaned our rooms or organized our thoughts. It has been said that an animator is first an actor, then a director, then a pencil pusher. The only thing that changes with the computer is to add keyboard tapper or mouse pusher to the third position with pencil pusher, where also sits the wonderful life of claymation and many other methods of animating.

Learning how to animate believable characters is not learning how to use the computer or learning how to draw or learning how to move a puppet. It is why and how the character moves that must come from the animator. It is going back to the cave paintings of animation. It is going back to the learning sessions of 1937 at Disney (not as Disney, but as the place where "Animating Force vs. Form" was the title of a discussion for the evolution of animation).

In "CPU," our character steps away from the means, the computer, to plan the "what and how," which is what made making "CPU" so much fun.

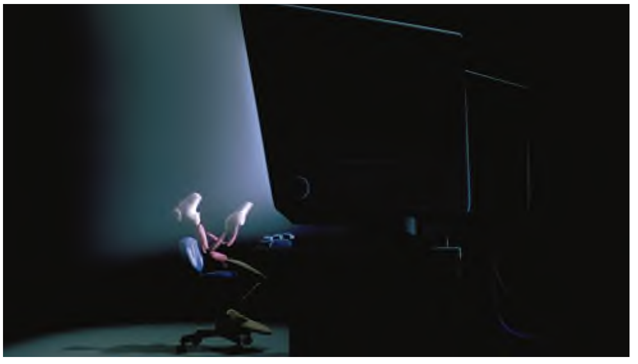

Loose and Sketchy Animation

Cassidy J. Curtis
University of Washington
Box 352350
Seattle, Washgton 98195 USA
cassidy@cs.washington.ed

In narrative character animation, engaging the viewer and conveying emotion and personality are of the highest importance. Early in the development of the medium, animators learned that an effective way to achieve these goals was to "present a unified single idea with nothing complicated, extraneous, or contradictory in its makeup."[1] But this principle is seldom applied successfully in computer animation because of the demands of photorealistic rendering.

It's What You Leave Out That Counts

Photorealism, like pornography, leaves nothing to the imagination. It presents the viewer with a world of objects complete with volume and texture, which is far more information than the viewer needs to get the point. Furthermore, unless great effort is devoted to every detail of modeling, shading, and lighting, much of that information will actually contradict the central idea, distracting the viewer. The need to cancel out these spurious impressions places an unreasonable demand on the animator.

The Power of the Sketch

A hand-drawn sketch can convey a lot of information in a much simpler package, if one draws only what is necessary. The gesture of line can also convey the energy of the artist's hand, or a character's mental state. And because a sketch is clearly an artist's interpretation of the world, it encourages the viewer to complete the picture by imagining the details that are missing. This engages the viewer's mind in a way that a photorealistic image cannot.

A filter for converting a 3D animation into a sequence of gestural sketches promises to be a powerful tool for motion testing, or for creating finished animations in a non-photorealistic style. This sketch describes one approach to implementing such a filter.

Technique

The "loose and sketchy" filter automatically draws the visible silhouette edges of a 3D model using image processing and a stochastic, physically based particle system. For input, it requires only a depth map of the model (Fig. 1) and a few simple parameters set by the user. First, the depth map is converted into two images:

- The "template" image (Fig. 2), in which each pixel represents the amount of ink needed in its immediate neighborhood. This image is obtained by calculating the magnitude of the gradient of the depth map, thresholding to give binary values, and then blurring the result.

- The force field (Fig. 3), a vector field that pushes particles along the silhouette edges. Such a field can be obtained by calculating unit vectors perpendicular to the depth map's gradient.

Next, particles are generated, one at a time, for a fixed number of particles. Each particle's initial position is chosen at random from within the template image, with a bias toward areas that need more ink. (No particles are ever born in areas that need no ink.) Acceleration at each timestep is based on the force field, with additional coefficients for randomness and drag. The particle is rendered onto the canvas as an antialiased line segment. If a particle wanders into an area that needs no more ink, it dies, and a new one is born in another random place. The particle also erases the dark pixels from the template image as it travels, so that those edges will not be drawn again. (This aspect of the technique is similar to previous approaches to pen-and-ink rendering.[2,3]

Results

By tweaking drag and randomness, a continuous gamut of styles can be generated, ranging from a tightly controlled technical drawing (Fig. 4) to a loose and gestural sketch (Fig. 6). The higher the drag, the tighter the style. These parameters also affect the character of motion, which can range from near-perfect smoothness to a lively staccato. The looser styles are particularly appropriate for animated characters, giving them a bit of life even when they stand perfectly still. It is also worth noting that for animation, it is not necessary to exactly render all of the silhouette edges all of the time. An error in a single frame will usually go unnoticed, because each frame lasts only a fraction of a second.

References

1 Frank Thomas and Ollie Johnston. Disney animation: The illusion of life. Abbeville Press, New York, 1981.

2 Michael P. Salisbury, Michael T. Wong, John F. Hughes, and David H. Salesin. Orientable textures for image-based pen-and-ink illustration. Proceedings of SIGGRAPH 97.

3 Greg Turk and David Banks. Image-guided streamline placement. Proceedings of SIGGRAPH 96.

Figure 1

Figure 2

Figure 3

Figure 4

Figure 5

Figure 6

Matte Painting in the Digital Age

Craig Barron
Matte World Digital
24 Digital Drive, #6
Novato, California 94949 USA
cbarron@matteworld.com
www.matteworld.com

Session: Invisible Effects

Since the turn of the century, matte paintings have been a mainstay in the filmmaker's repertoire of visual effects. These intricate landscapes, rendered on glass by talented artists, have long been used by filmmakers to create realistic illusions without exceeding the limits of a reasonable budget.

But today, many are beginning to question whether this technique is still a viable option in the digital age. Like many other filmmaking techniques, traditional matte painting is being completely redefined by computer-generated imagery. Traditional matte painting is quickly becoming a lost art, eclipsed by digital backlots and virtual sets.

Today's new digital capabilities vastly expand the ability of filmmakers to create realistic landscapes that do not exist in the real world. They can now quickly and easily darken a sky, add or remove a mountain, or test multiple color combinations. Before the digital era, it took weeks to re-paint a scene for the same effect. Digital technologies can enhance locations and sets in ways that would be too costly to construct by conventional means.

Digital techniques also free filmmakers from the two-dimensional limitations of traditional matte paintings. Now, for the first time, 3D rendering techniques allow the filmmaker complete freedom of camera movement through digitally generated environments.

But while new digital capabilities offer many exciting possibilities, they can never replace the artist. The tools may have changed from paint brushes to graphics tablets, but the artist's ability to visualize and mimic reality is still in high demand as a way to offset the enormous costs required to produce films today.

In fact, some traditional techniques often go unused even when the offer the possibility of superior results because the next generation of artists lacks the skills and training to use them. And the digital tools that artists need to paint on the computer today are still relatively crude when compared to traditional paints and brushes.

While working on effects for "Titanic," our studio had less than two weeks to complete a shot of the rescue ship Carpathia. It would have taken too long to render a 3D model of the ship, but one of our artists was able to paint it in less than a week. That traditional matte painting was then composited digitally into the scene.

Digital techniques also have the potential to save filmmakers time and money, but this is not always the case in practice. The time saved can easily be taken up with experimentation and comparison that was impossible with traditional techniques. And the cost of keeping up with the latest technology can make computer effects even more expensive than traditional ones.

Computer-generated realism may also not be right for every film. Traditional matte paintings, where artists create their interpretations of a scene instead of enhancing and combining photographic images, have a fundamental emotional quality that is still difficult to replicate digitally.

But as digital solutions improve, I believe traditional techniques will continue to shrink further and further into the background. In the future, I do not doubt that it will be routine to have whole sequences made in 3D environments. The number of films that use some kind of computer-generated special effect is growing at an astonishing rate, more than doubling every year.

Practitioners of the matte painting craft face many challenges and exciting possibilities as they make the transition into the digital age. For now, they must work in the gray area between traditional and digital techniques, and strive to combine the best of both worlds.

Hard Rain: A Journey from Title to Story

Brad Kuehn
Cinesite Digital Studios
1017 North Las Palmas Avenue
Los Angeles California 90038 USA
ill@cinesite.com

Hard Rain Credits
Director: Mikael Salomon
Producer: Cinesite Digital Studios
Visual Effects Supervisor: Brad Kuehn
Visual Effects Producer: Aaron Dem

At 4,997 frames in length, the opening sequence from "Hard Rain" shows how computer graphics can add possibilities to film and jettison you into a story, leaving you clueless as to where the real ends and the fake begins. For this sequence, high-end compositing was used to integrate 3D computer-generated water, rain, and fog (created in Arete, Dynamation, and proprietary code), farm equipment, and birds with live-action film shots, miniatures, matte paintings, and sky replacements, each with their own camera moves, to create one fluid journey into the film.

The direction of this opening epic started with the script, which very clearly stated that the journey was to take the viewer from the Paramount logo into a town endangered by the elements, ending on a closeup of the film's protagonist, Sherrif Randy Quaid. At the beginning of the production process, we determined what the shot would feel like, which elements would be of what type (live action, miniature, or computer-generated), and how they would bring the shot together. Because they were unable to find the perfect location, the team determined that the dam, levy, and approach to the town would all be miniatures. In addition, they filmed several live-action pieces: cars on highways, cattle moving onto dry land, and the final journey through the town. They also determined that computer-generated imagery and matte paintings would be used to augment backgrounds, skies, water, rain, and fog.

The Practical Shot
The bulk of the sequence was a cable-cam shot, which required the crew to set a 200-foot pole as far back as possible at the edge of town. The cable cam (a spacecam camera mount attached to a cable cam) had to stop just prior to crashing into (and severely injuring) the film's protagonist. Consequently, rather than stopping on a dime, the shot ended with an overriding zoom to the close-up on Quaid.

The Opening Shot
Starting with the Paramount logo, we wanted to give the viewer the feeling of the ominous nature of the situation. Cinesite's matte painters built several paintings to gradually replace the clouds that were included in the real logo, so one layer was added at a time until the final matte painting showed the sky engulfed in clouds. Another matte painting was used on the boom down from the logo to the bottom of the mountain.

The Dam Shot
Because the director wasn't entirely happy with the scale of the water and felt the viewer couldn't really feel the scale of the dam, Cinesite's artists augmented the model. They added CG water, created as a particle system in Dynamation, to the existing water coming out of the dam, making it appear more threatening and ominous.

We used CG water throughout for two reasons: because we couldn't shoot real water to go into the extreme shots used for the miniature plates, and because of the varying camera moves from one type of element to the next – whether live action to miniature, or miniature to miniature. CG water was used to smooth the transitions (hide the seams) between the sections.

CG Mist and Fog
Computer-generated fog was the miracle cure we used to hide, mask, and augment the journey throughout this sequence. We added fog during the shot as we pulled back from the base of the Paramount mountain and headed toward the dam. Fog was also used as a layer to smooth the transition between the rushing water from the dam and the water at the base of the dam, and as the camera travels over the final miniature town and into the real town.

CG Rain
Just as the ship is the star of "Titanic," you could almost say that CG rain was the lead in this piece. It started pouring at the beginning of the shot and continued through to the town, where the CG rain was replaced by real water from the rainbars.

Rotoscoping
At its simplest, rotoscoping is the process of tracing around a foreground element for the purpose of creating mattes or effects either around or behind it. It gets really complicated when you look at what our artists are tracing. We rotoscoped everywhere, especially the pullback from the dam, around the lamp posts and the bridge (frame by frame). We used additional rotoscoping wherever CG water met the miniature plate, and where buildings met the sky, we used rotoscoping to allow for inclusion of the sky-replacement matte painting.

Color Correction and Integration
Color correction is the art of making everything look exactly as it should. But you have to consider what looks right, and who determines that, and how much blue there should be in those very ominous clouds. Each live-action plate and miniature was shot under completely different lighting conditions, so they each had to be manipulated to match each other and the integrated CG elements. In addition, we had to turn the blue water to murky brown. Live action plates were shot on sunny days, and we had to make them overcast And all CG elements had to match to the miniatures and the live action plates.

Other Stuff
To make the final town shot work, we had to remove the cranes, the rainbars, and the cable on which the cablecam was attached. We had to add streetlights and cables back to the frame above Randy Quaid. And we had to take out all the camera equipment that was on the street beyond the car. We added a CG barn, silo, and trees because the scene takes place in Indiana, and we thought it would look cool to have a submerged silo. We added CG birds to bring a little touch of life as our camera flew over the town, but then we had to train them to fly where we wanted them to appear. We left the hose on the sidewalk because it didn't bother the director. Besides it was more cost effective that way.

Computing Procedural Soundtracks from Animation Data

Robin Bargar, Alex Betts, Insook Choi
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
Urbana, Illinois 61801 USA
rbargar@ncsa.uiuc.edu

This sketch describes a project to provide synchronous data-driven sound for animations in real-time and non-real-time display environments. The underlying objective is the capability for animation data to be applied to automated generation of soundtrack accompaniment. Sound production methods applied in this study include synchronous scheduling and reproduction of digitized sounds stored on computer disk, and real-time synthesis methods including wavetable synthesis, modal synthesis, and physically based models.

Sound is generated in correspondence to the data and dynamics of computational models. This correspondence can be reflected in the internal structure of a computed sound. The dynamics of a sound can be parameterized in synthesis algorithms, and these algorithms can be controlled by applying techniques for organizing and displaying numerical data. Where a graphical display is dynamically computed to visualize data or a numerical model, sound can be computed to respond to the same dynamics, to achieve a coupling of sound and image that is perceptually valid and informative about the states of the underlying computational model.

To create a software environment to support sound computation, a prototype of a general sound rendering system was developed. The prototype is capable of signal quality comparable to professional audio recording standards. The current implementation involves a software-only approach for UNIX platforms. Off-the-shelf peripherals and stand-alone audio hardware devices tend to be designed for music-industry applications and are not fully programmable nor controllable to an arbitrary number of parallel control parameters.

A research project has been underway to provide sound synthesis in software performed on a workstation CPU and output to a DAC in real time. The NCSA Sound Server (www.ncsa.uiuc.edu/~audio) was developed for sound computation and synchronization research, and for integrated audio-visual display. The Sound Server is a software-based sound production environment used by developers collaborating in VR and computer graphics research. The Sound Server demonstrates frame-accurate synchronization with variable-framerate real-time processes such as virtual-reality-based animations. This is accomplished in a client-server framework by computing sound as a parallel process with an internal time mechanism independent of processes such as graphics, dynamics, and drivers for control devices.

An interface protocol provides parameterized access to the dimensions of synthesis algorithms, and an interconnect mechanism encapsulates this dimensionality to support high-level sound organization. At a high level, relationships among sounds may be designated corresponding to conditions in an accompanying client application. We refer to specification of these high-level relationships as "sound authoring." A sound authoring system allows rapid prototyping of sound environments by mapping control dynamics extracted from independent computational models. These computational models may be found in computer graphic displays or other interactive applications.

This sketch describes the sound authoring framework and discusses methods for generating synchronous sound with computer animations and virtual environments. Synchronization is defined differently in real-time and non-real-time displays. In real time, synchronization may be approached as a range of temporal correspondences observed in parallel processes, rather than alignment of discrete events at a unity in time. A sound authoring specification is described as a mechanism performing automated observation of parallel processes in order to generate signals to be described as synchronous. Examples include a data-driven audio plug-in for 3D graphics modeling and animation environments, and live performance of virtual reality presented in a theatrical concert.

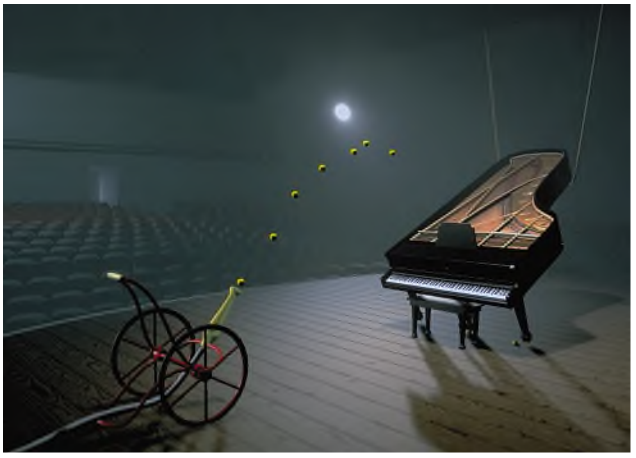

ImageTimer: A Traditional Approach to 3D Character Animation

Michael Blum, Nhi Casey, Greg Heflin
Walt Disney Feature Animation
500 South Buena Vista Street
Burbank, California 91521 USA
mikeb@fa.disney.com
nhic@fa.disney.com
gregh@fa.disney.com

This sketch describes a novel combination of tools that allow a computer animator to change the timing of key frames in an animation package by adjusting the spacing and order of images recorded at those key frames. The animator interactively manipulates the images in a frame-accurate playback device, flipping through unlimited timing tests and adjusting key frames at will.

Background

The computer graphics community expends a great deal of energy each year developing new techniques for artists to create ever-more-realistic models and animation. However, surprisingly little work has been done in creating animation tools that improve animator workflow – simplifying the process while allowing for the fine-grain control that film-quality animation demands. In addition, while many animation tools are designed with technically trained animators in mind, few tools attempt to bridge the gap between these 3D-savvy animators and the traditionally trained 2D animator. By excluding this latter group, we drastically reduce the available talent pool. Our work takes a step toward simplifying the animation process while introducing paradigms familiar to the traditionally trained artist. Note that by simplifying the process we do not restrict tool use. An animator is free to use the ImageTimer suite in conjunction with the pallette of tools provided by a third-party animation package.

This technique is best illustrated through an example:

Workflow

The animator begins by setting up poses for a character using the palette of tools provided by a modeling and animation package. At this point, no care is taken to time the poses; they can be arranged at consecutive frames. Next, a custom animation plug-in records snapshots of the scene from the frame buffer for each pose and sends these images to our proprietary, real-time, 2D, playback tool. Our playback system contains a host of animator-friendly features, including the ability to synchronize loops of images to an audio track, adjust the playback rate, and scrub through the loaded images. Most importantly, it contains an exposure sheet editor that the animator uses to interactively adjust the timing of the loaded images. The animator manipulates the images in the exposure sheet editor, viewing changes in our playback tool, until the desired timing is achieved. The animator then applies the timing changes back to the 3D poses by selecting the model hierarchy to be adjusted and clicking a button on our plug-in panel; all key frames for the selected nodes will be moved to match the timing created in the exposure sheet editor. Typically, an animator will use this technique to first construct a rough pose test by applying timing changes to the top-most node in the character's hierarchy. An animator can then add secondary motion by applying timing changes to sub-hierarchies of the model. Historically, these types of timing changes have been applied by hand using a 2D graph editor, a time-consuming, error-prone process that is unintuitive for most traditional animators.

The process described above directly correlates to the methodology our traditional artists use when drawing with pencil and paper. First, they draw key poses on consecutive pieces of paper. If, after flipping through these sheets of paper by hand, it is apparent that the animation is working, the animator can scan each drawing and use our proprietary playback tool to adjust timing. This animation is called a pose test. Once satisfied, the animator uses the information in the digital exposure sheet to create hand-drawn timing charts that instruct a junior animator how in-betweens should be drawn between the key poses. At first blush, this last step seems unnecessary in the digital world, where the computer performs in-betweening. However, timing-charts provide an efficient control mechanism that animators desire. Encapsulating the notion of a timing-chart and integrating it into the ImageTimer system is a topic for future work.

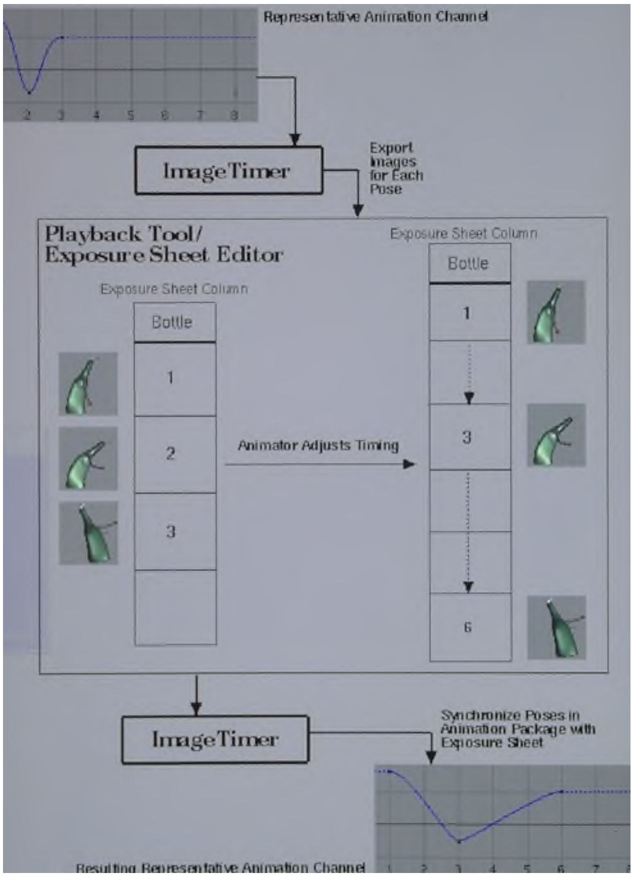

Session: It's Character, Stupid!

animations

The PDI Facial Animation System for ANTZ

Beth Hofer
Pacific Data Images
3101 Park Boulevard
Palo Alto, California 94306 USA
beth@pdi.com

In 1996, Pacific Data Images (PDI) and DreamWorks SKG entered into a co-production agreement to create "ANTZ," a fully computer-animated feature film. Among the many technical challenges we encountered was facial animation. This sketch includes a brief description of the facial animation issues encountered on the "ANTZ" project and the technical solutions we developed to meet these challenges.

Facial expression is used as a primary way of conveying a film's message or mood. In animated features, facial animation takes on added importance, because an inanimate object or an animal that the audience doesn't usually think of as having feelings and emotions can be the star of the film. The character must convincingly exhibit a range of expressions (a bashful smile, a quizzically arched brow, a dismissive sneer) that make the audience believe it is a living being filled with thoughts and emotions.

The facial animation system developed for "ANTZ" entails a complete set of muscles corresponding to the actual muscles of the human face. The combinations of muscles that create specific expressions in the human face are used in the same way to create expressions in the faces of the ants. For example, the frontalis and corrugator muscles are used in creating a sad expression. As the facial muscles are activated, they cause the face model to change shape in a way that respects the imagined physical characteristics of the material that would make up the face if it were real. Hard areas that correspond to bone change shape only slightly, but soft areas such as cheeks bulge out or sink in as they are compressed or stretched.

The facial system's interface is designed to allow an efficient, layered approach to animation. There are over 300 control elements in the characters' faces, consisting of muscles, bone movements, eye rotations, etc. Although they can be animated individually, they rarely are. Rather, they are combined into simple controls like "r_brow_mad", "l_eye_widen", or "l_mouth_sneer", each of which may activate several muscles in various amounts. These simple controls can be further combined into preprogrammed library shapes such as phonemes. The animator can quickly select desired library shapes and place them on the appropriate frames. The shapes are converted to the control combinations that make them up, which the animators can then individually refine.

It would be a great oversight not to mention the very talented PDI animators who have utilized the facial system to produce the unprecedented levels of acting in "ANTZ." The facial animation system has allowed them to animate powerful, emotional scenes. The characters have come alive and taken on a personality of their own. They interact with each other with knowing glances or sly smiles. They care for each other or show great disdain for one another, due in part to the wonderfully rich set of facial animation controls provided to the animators.

The StormRiders Feature Film

Mike Wong
Centro Digital Pictures LTD
601 HKITC
72 Tat Chee Avenue
Kowloon Tong, Hong Kong
mikool@centro.com.hk

Hong Kong filmmakers have produced many well-known modern action movies that have received very positive feed-back from worldwide audiences, and in the summer of 1998, Hong Kong's first heavy-duty visual effects movie premiered: "The StormRiders," based on Hong Kong's best-selling superhero comic book.

The story revolves around the feud between two young disci-ples and their mentor in ancient China. he movie contains over 500 visual effects shots, a record for Hong Kong movies. Digital effects employed in "The StormRiders" can be categorized into four types:

1 Digital Sets

There were many sets in the script that were difficult or very expensive to build or shoot, so we chose to build them in the computer. One of them is a Buddha in Rock that is over 20 stories high. The entire 3D model was built from scratch and 100-percent 3D-painted. The Buddha scene has a polygon count up to 4,000,000 and over 700 Mbytes of surface tex-ture maps. MetaNURBS (LightWave) and traditional patch modelling (3D Studio MAX) techniques were employed in the modeling process. A similar set, a fortress, also required large-scale modelling and texture mapping.

2 Digital Creature

A Kirin-in-Fire (Kirin is a Chinese mythical creature) played an important role in the movie, where it appears for the first time on-screen, in action, and, more importantly, in fire.

Since the comics have already given the audience a deep impression of the Kirin, the design and modeling process demanded a natural transformation of the 2D comic Kirin into a fully animatable 3D Kirin. Instead of using digitizing tools for the 3D modelling, we chose to build the Kirin without any digitized data and go straight into the 3D NURBS modeler. We wanted the model to be as animatable as possible, so in every modelling step, we tested the animation freedom of each component in parallel. Once the model was built, we no longer needed to go back to the modeling process again just because of some minor animation constraints. More impor-tantly, during the modelling process, we could demonstrate the animation freedom of the Kirin to the client so they could provide feedback well before the animation process. Several simple sculptures were also made just for reference and dis-cussion purposes.

This creature also demanded very accurate character anima-tion and sophisticated particle animation for the fire on its body. The whole Kirin is 100-percent-NURBS-modeled in Alias PowerAnimator, and all animations were completed using Softimage. Custom software links were written to allow exchange of static and dynamic skeleton information between the two software packages. All animations were then sent to our mini-RenderMan farm for Kirin body rendering.

The fire on the Kirin's body was animated using Wavefront Advanced Visualizer and Dynamation. This process required the Kirin to be polygonized and be sent into Advanced Visualizer for particle emitter clusters assignment. Dynamation is the only particle tool we used for the Kirin's fire, and we dedicated two Onyxs equipped with Reality Engine II to the particle animators for these fire animations.

Compositions were completed on several different platforms, but all employed the same standard composition require-ments (heat ripples and halos, etc).

3 Supernatural Phenomena

In many of the scenes, the leading characters use supernat-ural powers as weapons during battle. This inevitably meant that digital visuals had to be deployed in order to visualize the mighty power of these phenomena. An interesting use of CGI in depicting such power, was the small water tornado activat-ed by Fung (one of the superheroes) to attack the Kirin. We used traditional 3D animation to
produce the core of the tornado with a liquid look and then overlaid tons of particle onto it as splashes.

4 Digital Matte Paintings

There are also several dozen shots that employed digital matte painting.

Compared to the average cost and time spent in Hong Kong film productions, The StormRiders is an exception. The whole digital effects production employed around 20 animators, two matte painters, six compositing artists, and a small team of technical support over a period of more than 16 months. Equipment included over 50 processors, two Quantel Dominos, and over 400GB of hard-disk space.

Flubber

Tom Bertino
Industrial Light & Magic

When the professor's away, the Flubber will play. But for those who are familiar with Disney's original "Absent Minded Professor:" be warned. This is not your father's Flubber! This is living, breathing, thinking Flubber. Flubber with an attitude ... and rhythm! A dancin' Flubber that can subdivide so that it not only does a star turn, but also provides its own supporting chorus.

The character effects in "Flubber" presented Industrial Light & Magic with a new and unique set of problems, which called for equally new and unique solutions. How do you articulate a character made of slimy, sticky goo? How do you express acting with a character that has no face and rarely has definable body parts of any kind? How does light reflect and refract through a transparent character? These are a few of the challenges we faced, and we had to find answers that addressed aesthetic, performance, and technical concerns simultaneously. Form had to follow function, and this plunged us headlong into the realm of the unknown and unproven.

Previous CGI creatures tackled at ILM, such as elephants or dinosaurs, have had dependable physical structures. The various parts of the model may move, but that basic shape and structure, and the parts' relationship to each other, remain basically the same. Not so with Flubber. Flubber was constantly shape-shifting, undulating in a way that's appropriate for a semiliquid substance. It might pop out legs, or a head, or the whole creature could become one big mouth. In short, nothing was for sure with Flubber, or remained the same for very long. It was decided that such a creature called for Softimage's Metaclay program. This process creates a surface over animating blobby shapes. To our knowledge, this was the first time Metaclay was used to create a major character performance in a theatrical motion picture. Our classically trained animators were quite successful with this technology, since it resembles traditional drawn animation in the sense that you are redefining the form of your character frame-by-frame. The trick, though, is that Metaclay was never designed as a character acting tool, and it has a bit of an unpredictable nature, a mind of its own, not unlike the Flubber character! This necessitated more give and take between machine and operator than usual, but when all the factors were working together, it produced some marvelous results.

We wanted Flubber to look as real as possible, and to integrate well into the real-life environments in terms of lighting, shading, and texture. It was decided that by using a raytracer we could best represent Flubber's solid, glossy surface finish, show his hundreds of internal bubbles, and accurately reflect and refract his environment, his own appendages, and other Flubbers. In addition, it would allow us to see Flubber reflected in the elements of the set. This was the first show at ILM to use a raytracer for the main character, and the first to use one so extensively, since it was involved in every shot in which Flubber appeared. We developed surface and volume shaders for the Flubber raytracer, and also a technique to create "caustics" from light passing through the Flubber and illuminating the surface on which he was resting. In addition, it was necessary to find innovative ways to break down the scenes in the dance that contained over a hundred Flubbers so they would be more manageable, from both an animation and a rendering point of view. Not only were the Flubbers, their shadows, and their refractions split into separate passes, but each of these subdivisions had to be broken down into still smaller pieces, due to the sheer weight of information these shots contained.

ILM's work with the title character in "Flubber" was definitely something new and different in the realm of computer animation for theatrical feature films. It was a challenging and sometimes daunting journey, but along with that came a wonderful feeling of discovery that was shared by all the crew and, we hope, the audience as well.

SIGGRAPH TV

SIGGRAPH TV Film and Video

Introduction

SIGGRAPH Online / SIGGRAPH TV

For SIGGRAPH 98, SIGGRAPH TV and SIGGRAPH Online joined forces to create a prototype 21st century "all-media" production facility. This state-of-the-art television studio, production facility, and multimedia broadcast center was built entirely inside the Orange County Convention Center. The facility was used throughout the conference week to create and broadcast almost 60 hours of television programming and Internet content.

The facility occupied more than 6,000 square feet and included the very latest in television and Internet hardware and software. The programming created in the facility was broadcast to over 30,000 attendees during SIGGRAPH 98, provided locally to SIGGRAPH 98 shuttle buses, and delivered to 200 million Internet users in over 200 countries through a variety of "streaming" media and Web-based reporting.

A staff of more than 60 producers, editors, scriptwriters, graphic artists, computer administrators, online publishers, and technical personnel were brought together to produce this week-long event. The content included the world's best computer animations, interviews, and daily in-depth coverage of the 25th international conference on computer graphics and interactive techniques.

Designed as a showcase of technology, this facility was open to all SIGGRAPH 98 attendees. During production, attendees could take a self-guided tour through the facility and watch as television content was edited for broadcast and the Internet, and view first-hand the equipment used to create the content and manage the process.

SIGGRAPH Online Committee

Chair
Omar Ahmad
Netscape Communications

Steve Allison-Bunnell
Natureboy Media

Paul Hart
C/Net

Kevin Lahey
NASA Ames Research Center

Mario A. Jimenez
Orlando Sentinel

Mark C. Kilby
Orlando Sentinel

John Hulson
Discovery Channel Online

Meryle Mishkin
Netwizards

Danialle Weaver
Business / Technology Journalist

SIGGRAPH TV Committee

Chair
Dave Tubbs
Evans & Sutherland Computer Corporation

Jeanie Taus
Independent Consultant

Michael Pavlinch
imageDESIGN Videographics

Marc Parrish
Middle Tennessee State University

Lynn Finch
Finch Interactive Group

SIGGRAPH TV Global Post

SIGGRAPH 98 added a new component to SIGGRAPH TV: creation of a "virtual" production and animation facility. For the first time, individuals or organizations of any size could participate in creation of content for the SIGGRAPH conference by generating animations and video production work from any point on the globe.

MOdi Studio
2814 Madison Avenue
San Diego, California 92116 USA
+1.619.563.4819

Pixel Factory
4081-C L.B. Mcleod Road
Orlando, Florida 32811 USA
+1.407.839.1222

Avid Neo Geo
4390A 35th Street
Orlando, Florida 32811 USA
+1.407.423.9535

Age of Convergence

Nanotechnology, where the worlds of the real and the virtual converge. Personal applications of nanotechnology in daily life.

Nanotech Visionary: Charles Ostman
Director: Michael O'Neill
Producer: Kevin Cain
Lead Animator and Editor:
Ralph Fairweather,
Academy of Art College
Animators: Stephanie Alden,
Lewis Alexander, Edgar Arenas,
Katie Cole, John Gomes,
Matt Intrieri, Jordan Itkowitz,
David Kimelman, Sean McKenna,
Alan Quadros, Andy Schlussel,
Michael Shahan, Brad Steinberg,
and Amanda Webb
Narrator: Maggie Ziomek
Actress and Writer: Sabrina O'Neill
Music: Mahal
Technology Partners: Art+Com
(Berlin), Biovision (San Francisco),
Cyberware (Monterey),
Parker Instruments (San Francisco),
Digibotics (Austin)

Contact
Michael O'Neill
1215 Bay Street, #7
San Francisco, California 94123
USA

Adam & Eve

Julie Janower directed, wrote, animated, and produced this short computer-animated film about the destruction of the world.

Contact
Julie Janower
1133 10th Street, #104
Santa Monica, California 90403
USA
+1.310.458.3228

Adventures of Spiderman

Stereoscopic CG Spiderman battles Dr. Octopus, Electro, Hydroman, and other super-villains in the first combination of motion-base and rollercoaster technology. Using 26 8 per/70mm projectors, all CG environments are blended with practical set pieces to transport the viewer inside Spiderman's world.

Producers: Jeff Kleiser and
Diana Walczak

Kleiser-Walczak Construction Co.
Line Producer: Patrick Mooney
Line Producer (Test Period):
Erika Walczak
Technical Supervisor:
Jeffrey A. Williams
Head of Software: Frank Vitz
Lighting Director: Kathy White
Animation Director: Larry Weinberg
Modeling Supervisor: Scott Palleiko
Character Animators: Derald Hunt,
Dana Peters, Fabio Tovar,
Phearuth Tuy
FX Supervisor: Greg Juby
Character Supervisor: Josh Reiss
Animators: Ed Batres,
Michael Clausen, Randy Goux,
Ray Haleblian, Beau Janzen,
Doug Kingsbury, Jeff Lew,
Mark Pompian, Carriann Powrozek,
Kody Sabourin, Paul Stocker
Artist Storyboard: Kent Mikalsen
Production Manager:
Tom Hendrickson
Digital Assistants: Billy Barnhart,
Robin Cookis, Lee Mylks,
Slavica Pandzic
Film Department Manager:
Martha Small
Systems Administrator: Joe Hall
Assistant Systems Administrator:
Scott Lord
Production Coordinators:
Arch Gibson, Santo C. Ragno
Office Manager: Andrea Cancilla
Accountant: Kim Alice
Production Assistants:
Scott Kirchner, Molly Windover
Office Assistant: Erick Ungewitter

Universal
Director: Scott Trowbridge
Production Designer: Thierry Coup
Art Direction: Phill Bloom, Eric Paar
Producer: Mark Rhodes
Project Coordinator: Greg MacLaurian

Contact
Jeff Kleiser
Kleiser-Walczak Construction Co.
87 Marshall Street
North Adams, Massachusetts
01247 USA
+1.413.664.7441
+1.413.664.7442 fax
jeff@kwcc.com

An Indirect Consciousness

A poetic piece that communicates a significant reason for the importance of intelligent life and all life on earth and on other planets, for we're all the conscious thinking elements of the entire universe. We interpret it through science, arts, and technology.

Online Editing: Jacky Kamhaji
On-Camera Work: Kenneth Eng

Contact
Danny Kamhaji
51 Monitor Street, #3
Brooklyn, New York 11222 USA
+1.718.389.8481
+1.718.389.8481 fax
friction@dcdu.com

and so, she departed...

This story is about the last minutes in the life of a little girl. A rag doll sits on a night table in her room, while sounds of the past bring back memories. The hardest thing about losing someone is keeping the good memories alive without forgetting the bad ones.

Director: Celeste Ramirez
Producer: Ringling School of
Art and Design
Concept and animation
Audio mix: Celeste Ramirez
Faculty Advisor: Ed Cheetham
Video/audio support: Phil Chiocchio
Software: Alias PowerAnimator 8.2,
Composer
Hardware: SGI O2

Contact
Celeste Ramirez
c/o S. Trovas
Ringling School of Art and Design
2700 N. Tamiami Trail
Sarasota, Florida 34234 USA
+1.941.359.7536
+1.941.359.7517 fax
strovas@rsad.edu

Bob & Scott

A television series in CGI cartoon mood, with key-frame animation.

Directors: Olivier Bonnet and
Laurent Bounoure
Producer: TF1 / Protecrea

Contact
Guillaume Hellouin
SPARX*
91, rue Lauriston
75016 Paris, France
+33.1.56.28.01.28
+33.1.56.28.04.29 fax
sparx@imaginet.fr

Beryllium

Producer/Audio: Todd Sines

Contact
Neal McDonald
ACCAD
The Ohio State University
1224 Kinnear Road
Columbus, Ohio 43212 USA
+1.614.292.1041
+1.614.292.7776 fax
mcdonald@cgrg.ohio-state.edu

Crash Mapping

The effect of a new technique called "crash mapping." Only from a general polyhedra geometry and a gray-scaled crash mapping image, the object is divided into pieces automatically. The shape of the pieces corresponds to the crash mapping image.

Arghyro Paouri and Christian Blonz

Contact
Youichi Horry
Central Research Lab.
Hitachi
1-280 Higashi-Koigakubo,
Kokubunji
Tokyo 185 Japan
+81.423.23.1111
+81.423.27.7754 fax
horry@crl.hitachi.co.jp

The Dragon

Student portfolio. Animation study of character animation and rendering.

Directors: Sebastien Bruneau,
Cedric Schmitt
Producer: ICARI Institute

Contact
Luc Larouche
ICARI Institute
85 St-Paul West, Suite 31
Montréal, Québec H4Y 3V4 Canada
+1.514.982.0922
+1.514.982.0288 fax
larouche@icari.com

Distortion

Death and life and the reason for existence in a distorted world.

Directors: Tshyoshi Mizoguchi and
M. Hirotsugu Motoyama
Animator: Jyan Kinoshita

Contact
Hiromi Habuto
2-3-4 Nagata Minami, Minami-Ku
Yokohamashi 232 Japan
+81.45.742.3050
habuto@interlink.or.jp

Equis

"Equis" is the Spanish pronunciation of the letter x. X and what this symbol actually resembles is the main topic throughout this visual exploration.

Hardware and software support:
School of Visual Arts

Contact
Fabian Tejada
43-59 161st Street, 3rd Floor
Flushing, New York 11358 USA
+1.718.899.8464
sheol7@rocketmail.com

The EyeCue System

What will educational technology look like in the future? How will advances in artifical intelligence, 3D graphics, and animation combine to make learning more effective and fun? The NCSU Intellimedia group introduces The EyeCue System, featuring agent EyeCue and avatar Wizlo who help guide students through vivid, immersive learning environments.

Mary Ashley, Wes Leonard,
Tim Buie, Michael Cuales,
Alexander Levy, Rob Gray,
Frank and Susan Wimmer

Contact
Patrick FitzGerald
Intellimedia
2208 Creston
Raleigh, North Carolina 27608 USA
+1.919.515.3655
+1.919.515.7896 fax
pat_fitzgerald@ncsu.edu

Falling Idol

In this short film based on old movie styles, the idol (Mars, the God of War) is falling to the ground.

Contact
Keiji Yamauchi
Dentsu Tec Inc.
Creative Headquarters
1-8-9 Tsukiji, Chuo-Ku
Tokyo 104 Japan
+81.3.5551.8828
+81.3.5551.9405 fax

Fool Running

Twelve characters are separated into four groups of three men. They are forced to run in four different stages.

Producer: Takashi Yakushiji
Animators: Masahumi and Yoshiyuki Matsukuma

Contact
Tomoyuki Harashima
2739-122 Negoya Tsukui-machi, Tsukui-gun
Kanagawa 220-0203 Japan
+81.427.84.3015
+81.427.84.6609 fax
tomoyuu@ca2.so-net.or.jp

Frankenmouse

Director: Conchita Labra
Producer: Universitat Illes Balears (Maisca)

Contact
Juan Montes de Oca
Universitat de les Illes Balears (Maisca)
Ctra. Valldemossa Km.7.5
07071 Palma de Mallorca, Spain - Baleares
+34.71.172674/172995
+34.71.173003 fax
info@studio1.vib.es

Gone Forever

Samantha Sharpe, Natasha Sharpe, and David Biedny

Contact
Stuart Sharpe
42 Sunset Way
San Raphael, California 94901 USA
+1.415.256.9872
+1.415.256.9759 fax
gone@ssharpe.com

The Great Paper Deadline

Comical silent-movie-style time lapsed during the final hours before the SIGGRAPH 98 paper submission deadline at the University of North Carolina.

Directors: Mary C. Whitton, Todd E. Gaul, Andrei State, Andrew Ade

Contact
Andrei State
University of North Carolina at Chapel Hill
CB 3175 Sitterson Hall
Chapel Hill, North Carolina 27599-3175 USA
+1.919.962.1810
+1.919.962.1799 fax
andrei@cs.unc.edu

H20

What happens when you leave the water running too long.

Idea, modeling, and secondary animation: Aaron J. Hartline
Lighting, Effects, and Compositing: Adam Holmes
Animation: Mike Laubach

Contact
Aaron J. Hartline
1737 Lake Street
Whiting, Indiana 46394 USA
+1.219.473.9116
AJHARTLINE@aol.com

Herrmann Hall

Herrmann Hall shows an interactive, real-time walkthrough of the Naval Postgraduate School's virtual administrative building. The Herrmann Hall model is blue-print and real-world accurate and was constructed in 400 hours using Multi-Gen. The walkthrough software is NPSNET-IV utilizing the University of Pennsylvania Jack ML model.

Director: Fred Zyda
Producer: Michael Zyda
Video Post-Production: Fred Zyda
3D Modeling: John Locke

Contact
Michael Zyda
Department of Computer Science
Naval Postgraduate School
Spanagel Hall 252, Code CS/Zk
Monterey, California 93943-5118 USA
+1.408.656.2305
+1.408.656.4083 fax
zyda@siggraph.org

Intruding

In this story, all the characters intrude in the real world.

Contact
Ming-Huei Shih
139-35 35th Avenue, Apt. 1D
Flushing, New York 11354 USA
+1.718.886.2513
+1.718.886.2513 fax
jackyshi@aol.com

Islands of Adventure Web Game Tour

University of Central Florida and Universal Studios

Contact
David Haxton
University of Central Florida
Orlando, Florida 32816-1324 USA
+1.407.823.3110
+1.407.823.6470 fax
haxtond@aol.com

Jurupa

It's night-noon in a deserted graveyard. Hip bones dance to the smoky beats of the Jurupa (Hurypa). O Beato Agitato!

Director/Producer: John Clark Matthews, Sony Pictures Imageworks

Contact
Sande Scoredos
Sony Pictures Imageworks
9050 West Washington Boulevard
Culver City, California 90232 USA
+1.310.840.8312
+1.310.840.8100 fax
sande@spimageworks.com

Kachina Doll

This animation gives spirit to wood-carved kachina dolls, brings them to life, and displays the vitality of the spirit of the Hopi Indian.

Contact
Vicky Yu-tzu Lin
510 East Oglethorpe Avenue
Savannah, Georgia 31404 USA
+1.912.898.8363
ylinvicky@yahoo.com

Knife Runner

Directors: Diana Velilla and Ruben Villoria
Producer: Universitat Illes Balears (Maisca)

Contact
Juan Montes de Oca
Universitat de les Illes Balears (Maisca)
Ctra. Valldemossa Km.7.5
07071 Palma de Mallorca, Spain - Baleares
+1.34.71.172674/172995
+1.34.71.173003 fax
info@studio1.vib.es

Labyrinth

The global society is becoming more and more complicated, just like a three-dimensional labyrinth.

Contact
Masa Inakage
2-24-7 Shichirigahama-Higashi
Kamakura, Kanagawa 248 Japan
+81.467.32.7941
+81.467.32.7943 fax
inakage@medi-studio.co.jp

Lotus Spring

Opening animation for the CD-ROM "Lotus Spring," which is in production by Xing Xing Computer Graphics Inc. The animation shows the scenery of the computerized 900-acre Garden of Perfect Brightness in Beijing, which was burned to the ground in 1860 during the Opium War.

David Botta, Chihan Chen, Paul Chin, and Jack Chu

Contact
Lifeng Wang
Xing Xing Computer Graphics Inc.
2366 Main Mall, Room 047
Vancouver, British Columbia
V6T 1Z4 Canada
+1.604.822.6994
+1.604.822.5485 fax
wang@cs.ubc.ca

Meta Baron

Creative test for a full-CGI feature film.

Director: Herve Masseron
Producer: SPARX* / Chaman

Contact
Guillaume Hellouin
SPARX*
91, rue Lauriston
75016 Paris, France
+33.1.56.28.01.28
+33.1.56.28.04.29 fax
sparx@imaginet.fr

NPSNET: An Amphibious Virtual Environment

This video shows the NPSNET-5 prototype as an amphibious virtual environment. The landing craft can be driven by a human operator and dock with the larger ship. The human can get off the landing craft and walk through the ship engineering model. The software is written with Coryphaeus Easy Scene for an SGI Reality Engine-2.

Director: Fred Zyda
Producers: Michael Zyda, John Falby
Video Post-Production: Fred Zyda
Software: Didier LeGoff and Ken Watsen

Contact
Michael Zyda
Department of Computer Science
Naval Postgraduate School
Spanagel Hall 252, Code CS/Zk
Monterey, California 93943-5118 USA
+1.408.656.2305
+1.408.656.4083 fax
zyda@siggraph.org

paysage sylvain

This work, based on Debussy's piano music "Deux Arabesques - 1ere," creates the atmosphere of French impressionism in computer animation.

Contact
Sachiho Murata
Sonology Department
Kunitachi College of Music
5-5-1 Kashiwa-cho, Tachikawa-shi
Tokyo 190-8520 Japan
+81.42.535.9562
+81.42.534.3717 fax
sachiho@kcm-sd.ac.jp

PMU Jouez avec vos émotions

Director: Henri Barges
Producers: Byzance Production and SPARX* / Editing Studio

Contact
Guillaume Hellouin
SPARX*
91, rue Lauriston
75016 Paris, France
+33.1.56.28.01.28
+33.1.56.28.04.29 fax
sparx@imaginet.fr

Project Wivern

A powerful and rhythmic animation.

Director/Producer: Toshiyuki Aoyama

Contact
Toshiyuji Aoyama
2-24-2 202 Kokubunzishi, Honmachi
Tokyo, Japan
+81.48.471.3574
wivern@ceres.dti.ne.jp

Rococo #506

Rococo #506 is an animated visual exploring of the role of the designer in the television commercial industry today. It is about the electronic synthesis of art and its purpose, the commerce of style for sale.

Music: Young American Primitive
Equipment: Windmill Lane Productions

Contact
Stefan Smith
Windmill Lane Productions
1558 10th Street
Santa Monica, California 90401
USA
+1.310.576.1344
+1.310.576.1505 fax
stefan@windmill-lane.com

Sesame Street Revisited

A film that recalls the energetic use of animation on children's TV.

Contact
Ralph Destefano
Electronic Visualization Laboratory
University of Illinois at Chicago
3610 North Keeler, #2
Chicago, Illinois 60641 USA
+1.312.902.8510
+1.312.902.8510 fax
ralph@evl.uic.edu

Test 001

A man is entering a laboratory to test a new video game. The theme is a very interactive fight game.

Producer: Université de Provence

Contact
Franÿois-Xavier Aubague
Université de Provence
13, les Grands Champs
63 360 St-Beauzire, France
+33.6.14.49.24.15

Tibetan Dreams

This visual essay is an experiment in the new use of production audio soundtracks to animate objects and effects in 3D space. Using Tibetan religious music as a metaphor, various audio amplitudes were extracted and converted into animation channels. Audio was then used to birth the particles, provide wind effects, and separately offset animation of both the individual curtains and sacred wall hangings. The drum segments were extracted to control the flame intensity, as well as to trigger the "mandela effect" at the climax of the piece.

Producer: Ivylee Lim, Robert Brodey, Gelek Gyaltong, Jeff Wagner, Leo Chan, Alex Wai, Rob Bairos, Side Effects Software, Andrew Alzner, Alex Orgill, the Tibetan Government in Exile

Contact
Michael Carter
Two Quacks and a Baboon
c/o Side Effects Software Inc.
477 Richmond Street West,
Suite 1001
Toronto, Ontario M5V 3E7 Canada
+1.416.504.9876
+1.416.504.6648 fax
mcarter@sidefx.com

To Begin With

This work was composed in an attempt to visualize a sacred mountain in Nara Prefecture, known to be the district where the erstwhile metropolis of Japan existed.

Producer: Yukiko Okamoto

Contact
Keisuke Imanishi
Tetra Vision
No. 1031-178, Higasi Tomigaoka
3-Chome, Nara-Shi, Japan
+81.742.41.1549
+81.742.46.6489 fax
imanishi@kcua.ac.jp

Trade Secrets of the Violin Masters

This technological fusion unites various motion capture, scanning, and animation techniques as part of an initiative to capture and replicate the playing characteristics of the world's master violinists both for preservation of their legacy and for study within an interactive educational program series.

Director: Laurence Leydier
Producers: IMIT, Interactive Media Productions, Nothern Digital, Cyberware, MIT, Biomechanics, Waxworks

Contact
Laurence Leydier
Innovative Music Instructional Technology
15 English Place
Winnipeg Manitoba R2M 5J1
Canada
+1.204.254.7857
+1.204.256.6733 fax
lleydier@total.net

Tune Quest

An animation created without motion capture.

Director/Producer:
Toshiyuki Aoyama

Contact
Toshiyuji Aoyama
2-24-2 202 Kokubunzishi,
Honmachi
Tokyo, Japan
+81.48.471.3574
wivern@ceres.dti.ne.jp

Variation from "Tchaikovsky Pas de Deux"

Alias | Wavefront Kinemation was employed to realize character movement in this work based on the "Tchaikovsky pas de Deux."

Contact
Mitsuyo Hashida
Sonology Department
Kunitachi College of Music
5-5-1 Kashimwa-cho, Tachikawa-shi
Tokyo 190-8520 Japan
+81.42.535.9562
+81.42.534.37?7 fax
mitsuyo@kcm-sd.ac.jp

The Wonder of it All

The oneness of discovering true love, and the joyful sensations brought about by the lovers' communication.

Director: David Dohan
Producer: Ringling School of
Art and Design
Concept, animation, audio mix:
David Dohan
Music Composer: Timb Kuder
Faculty Advisor: Claudia Cumbie-Jones
Video/audio support: Phil Chiocchio,
Sy West
Software: Alias PowerAnimator v8.2,
Composer
Hardware: SGI O2

Contact
David Dohan
c/o S. Trovas
Ringling School of Art and Design
2700 N. Tamiami Trail
Sarasota, Florida 34234 USA
+1.941.359.7536
+1.941.359.7517 fax
strovas@rsad.edu

ACM SIGGRAPH
1515 Broadway, 17th Floor
New York, New York 10036 USA
+1.212.626.0500
+1.212.764.5537 fax