

NEW ORLEANS
2000
SIGGRAPH

Conference Abstracts and Applications



Conference Abstracts and Applications

COMMUNITY OUTREACH

ADELE NEWTON
UNIVERSITY OF TORONTO

EDUCATORS PROGRAM

VALERIE A. MILLER
GEORGIA STATE UNIVERSITY

EMERGING TECHNOLOGIES

J. JEFFREY CLOSE
INDEPENDENT CONSULTANT

PANELS

ALIZA CORSON
WALT DISNEY FEATURE ANIMATION

SIGKIDS

MARC BARR
MIDDLE TENNESSEE STATE UNIVERSITY

SKETCHES & APPLICATIONS

TOM APPOLLONI
HARRIS CORPORATION

Conference Abstracts and Applications

COMPUTER GRAPHICS
Annual Conference Series, 2000

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

ISBN 1-58113-209-3
ISSN 1098-6138
ACM Order No. 435003

Additional copies may be ordered pre-paid from:
ACM Order Department
P.O. Box 12114
Church Street Station
New York, New York 10257 USA

Or, for information on accepted European currencies
and exchange rates, contact:
ACM European Service Center
108 Cowley Road
Oxford OX4 1JF
United Kingdom
+44.1.865.382338
+44.1.865.381338 fax

Credit card orders from U.S. and Canada:
800.342.6626

Credit card orders from the New York
metropolitan area and outside the U.S.:
+1.212.626.0500

Single copy orders placed by fax:
+1.212.944.1318

Credit card orders may also be placed by mail.

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

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

Copyright © 2000 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page or initial screen of the document. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept., ACM Inc., fax +1.212.869.0481, or permissions@acm.org.

COMMUNITY OUTREACH

4

EDUCATORS PROGRAM

18

EMERGING TECHNOLOGIES

70

PANELS

102

SIGKIDS

158

SKETCHES & APPLICATIONS

168



Contents

- 9 CyberInstitute Technology Magnet
School @ Riverside High School
- 10 If You Build It: The Story of an
Online Collaborative Environment
- 12 techARTS: Urban Girls Get A
Handle on Computer Graphics
- 16 TeleCommunity: Visible Concept

COMMUNITY OUTREACH

ADELE NEWTON
UNIVERSITY OF TORONTO



Committee

CHAIR

Adele Newton
University of Toronto

ADMINISTRATIVE ASSISTANT

Etta DiLeo
The Write Approach

COMMITTEE

Lisa Carroll

Gregory T. Dobson
University of New Orleans

Lorraine Flaherty

Mark Jones
Cyberstage Communications

Marty Sixkiller
Pacific Data Images

JURY

Adele Newton
University of Toronto

Lorraine Flaherty

Mark Jones
Cyberstage Communications

THE COMMUNITY OUTREACH PROGRAM was inaugurated at SIGGRAPH 96 in New Orleans to provide tangible benefits, beyond those of rudimentary economic value, to the local community that hosts the annual SIGGRAPH conference. When the conference returned to New Orleans in 2000, the Community Outreach program was extended to include activities that could provide ongoing benefit to the local New Orleans community.

Community Outreach is designed to provide information and inspiration for high school students and teachers in New Orleans and around the world. Through close collaboration with the SIGGRAPH Education Committee, the New Orleans Public School Board, the Louisiana Children's Museum, and the New Orleans Museum of Art, and the volunteer efforts of industry professionals, SIGGRAPH 2000 Community Outreach presented a series of pre-conferences seminars and hands-on workshops:

NOVEMBER, 1999

To provide an overview of the computer graphics industry, international industry representatives from three countries volunteered their time to speak to students from five New Orleans high schools in November, 1999. These presentations featured information about the history and current state of the industry, current and future career opportunities, and educational options available to qualify for those careers.

FEBRUARY, 2000

Marty Sixkiller from Pacific Data Images presented a two-hour professional development session on how art, mathematics, and technology were applied in production of the feature film "Antz." This session at the New Orleans Museum of Art was attended by New Orleans high school art, mathematics, and technology teachers.

JUNE, 2000

New Orleans high school teachers participated in Creating Characters, a one day workshop on art and technology taught by Michael Carter, a veteran computer animator and technical director. Presented at the New Orleans Museum of Art, Creating Characters provided teachers with hands-on experience in character development using both traditional methods and computer tools.

During the conference, SIGGRAPH 2000 Community Outreach activities included programs for young children and high school students:

CAMP SIGGRAPH

A five-day camp for children ranging in age from 9 to 12, selected by the New Orleans Public Schools from the inner city district. The camp at the Louisiana Children's Museum delivered hands-on instruction in topics including photography, drawing, storytelling, storyboarding, and Web design and creation. The camp curriculum was developed and delivered by SIGGRAPH volunteers.

SIGGRAPH HIGH

New Orleans high school teachers nominated students to attend SIGGRAPH High, which included portions of SIGGRAPH 2000 and special sessions developed specifically for the high school students. Fifty students participated in this program, which included presentations from industry representatives on computer art, emerging technologies, special effects in film, computer animation, scientific visualization, and careers in computer graphics.

SIGGRAPH 2000 HIGH SCHOOL STUDENT SCHOLARSHIPS

In conjunction with the SIGGRAPH 2000 sigKIDS program, this initiative provided full conference scholarships to 10 New Orleans high school students nominated by their teachers.

COMMUNITY OUTREACH PRESENTATIONS

This series of papers and panels on innovative community outreach programs (offered as part of the Educators Program) provided information and insight for teachers, administrators, and curriculum designers.

As the leading international association for computer graphics professionals, SIGGRAPH is uniquely positioned to develop and implement the SIGGRAPH 2000 Community Outreach program. The program has had profound effects on the local education community in New Orleans and, through the presentations in the Educators Program, has the potential to inspire innovative outreach programs internationally.

This multimedia presentation is a chronology of how a high-technology (digital-media-based) magnet school program was established in an at-risk high school in El Paso, Texas. With still and video images, the presentation showcases “how to do it” — how to prepare students for higher education and employment in the exploding computer graphics arena.

In the spring of 1992, a former high school art teacher and computer graphics artist (Dennis Crawley) and a computer literacy teacher (Diego Morales), both of whom were working as educational technologists in the central El Paso School District office, combined their talents and proposed establishment of the school district’s Digital Media Department. The department would be operated by high school students, who would be trained in digital media so they could support the school district with resources from graphics to video production to online training.

The premise for the department was five years ahead of its time. Year after year, the concept was brought back to the El Paso superintendent of schools. Finally, in the spring of 1997, the superintendent asked Diego Morales to establish a technology magnet school program at Riverside High School, Ysleta Independent School District, El Paso. Dennis Crawley was immediately named the program’s lead teacher and computer graphics expert. The result: an educational learning environment unlike any other in the El Paso area.

The technology magnet school program was created during the summer of 1997 and throughout the 1997-98 school year. Many steps were taken before the first students walked through the doors. First and foremost was deciding what the students would take with them when they completed four years of study at the technology magnet school. Of course, they would receive a high school diploma, but the academic component had to be at the highest possible level. Programs were designed to ensure that students achieved sufficient academic success for entrance into a four-year college or university. This requirement was met by requiring that all students graduate with Texas “Recognized” diploma standards. The next level up was the “Distinguished Achievement Plan.”

The second and parallel academic strand was the technology component. Here it was determined that the students should take courses that would familiarize them with the “Industry” (what real-world graphics work environments were using. Standardization requirements led to selection of Macintosh computers and supporting graphics software: Adobe’s After Effects, Illustrator, Photoshop, and Premier; Metacreation’s Infini-D, Logomotion, Bryce 3D, Poser, and Show; Appleworks; Microsoft Office 98; NetObject Fusion; Newtek Lightwave; Strata StudioPro; and Macromedia’s Authorware, Director Multimedia Studio, and Soundedit 16.

The classroom environment played a significant role in creation of a comfortable and highly motivating learning climate. So the school created a CyberCafe modeled on a 1950s diner. The room consists of 14 diner booths, six high tables with stools, marbled laminate tops with chrome edging and table bases, neon-trimmed windows, and, of course, a 50s juke box. Twenty-three networked PowerPCs, with software and video equipment, are spread throughout the room.

A marketing campaign began with creation of still images that would be shown in the trivia and advertising segments prior to movies in local theaters. Print advertisements were placed in local news publications. Two 30-second commercials were developed and aired for 12 weeks before the school year. Information sessions were presented to middle school students and parents. An information booth was set up at a local mall, and flyers were mailed out to every incoming freshman throughout the school district. As a result, 104 students were recruited for the technology magnet school’s first year (1998-99).

For the second year, recruitment was limited to 60 students. A new CyberStudio, an integrated computer classroom with a television production studio, was installed in March 2000.

The school’s curriculum is based on the graduation requirements for the State of Texas. Additional requirements were established by the school district. CyberInstitute students must enroll in a technology-based course in each of their four years. During the first year, students take multimedia development. Here they receive their foundation skills in 2D and 3D graphics and animation along with entry-level video production techniques. The second year offers more advanced programs and skills, and students produce a video magazine focused on technology. Years three and four will be spent on specific areas of student interest and internships and mentorships with local technology-based industries and agencies.

All the aspects of the highs and particularly the lows, problems, and politics of establishing a technology magnet school are shared in this presentation. It includes student work from the end of the school year and the latest reports on recruiting and program enhancement.



The Collaboratory Project at Northwestern University is establishing a network-based collaborative environment that involves education, research, cultural, not-for-profit, business, and industry communities in the greater Chicago area and beyond. It is supported through a grant from the Ameritech Foundation. Staff, space, and computing infrastructure are part of Northwestern University's Information Technology organization. The staff consists of seven people whose responsibilities range from programming, information architecture, and multimedia design to project support and training (in the field and in our computer classroom), planning, evaluation, and management.

The Collaboratory Project has evolved from an early focus on the off-the-shelf networking tools and training that were highly successful in building the university's networked environment to a custom-designed, Web-based environment where participants join various "communities" based on their fields of interest. Communities in this context are database-driven shell applications that generate HTML pages following templates designed for various information architectures and functions. Within the communities, participants can collaborate in a variety of ways, from sharing information through Web pages and participating in threaded discussions and chat sessions to searching for community members with common interests and proposing and developing collaborative projects. A key feature of the environment is its use of a database-driven infrastructure. This lets participants devote their time to providing content, while the server assembles and formats Web pages on demand. In addition, responsibility for creating projects and submitting and approving content rests with the participants through a multi-level scheme of distributed management.

Most Collaboratory Project participants are K-12 educators within the greater Chicago area, but the project also includes participants from the state of Illinois and even several international locations. Typically, the project works with small groups of teachers within a school to develop specific projects. Often these groups are part of larger district or even state-wide projects. The project also works with the not-for-profit sector, and Chicago-area museums and cultural institutions.

Supported projects include:

- The I-57 project gathered images and texts on local histories of towns along Interstate 57, which crosses Illinois from north to south.
- The Internet Book Club, where teachers and students share book reviews, essays, poems and stories. The club includes facilities for finding and contacting teachers with common interests.
- MICNet (developed with the Northwestern University School of Music), a community where students can post musical compositions and receive comments and critiques from a composer-in-residence, music education students at Northwestern, and other students, who can even download, edit, and repost compositions as part of their commentary.
- Survey Studio, where participants can create online surveys that generate online graphics and downloadable statistics on results.

- The Repository, where educators can create a hierarchical classification as an interface to a knowledge base, with custom templates for submission, management, and display. In its first deployment, the Repository is being used as a storehouse of science fair projects.

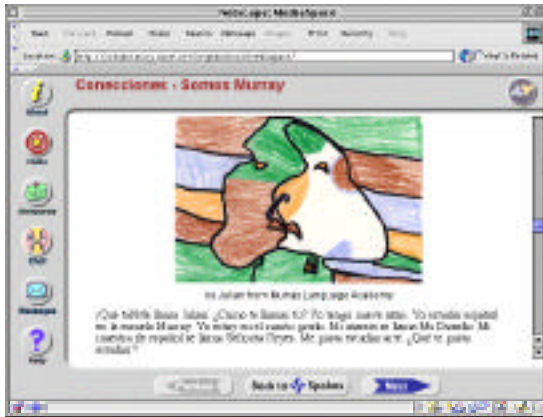
The project will further develop its collaborative environment during the coming year. A first step will be implementation of a portal site that will further integrate the different communities to which each participant belongs. Participants will gain immediate access to their assets within the various communities through a simple login procedure to a personal gateway. Similarly, projects that span several communities will be accessible through a single gateway. The portal design will further enhance opportunities for collaboration by providing each user with a profile that can be used to bring together participants with common interests and alert participants about new developments within their fields of interest. Design of the portal site is based on an object-oriented approach. If this approach proves as useful and flexible as expected, the object model will be extended to all of the communities.

Once you have built an environment where collaboration can take place, the critical question is: "Is it really working?" The Collaboratory Project evaluates results in several ways. Some of its Web pages compile statistics on projects and participants. For example, the K-12 Report page (collaboratory.nunet.net/cwebdocs/particip/participants/schools_report.cfm) lists some 111 schools and 11 multi-school projects. Forty-eight schools are in the city of Chicago, another 43 are spread among 15 suburban school districts, and the 20 remaining schools are located elsewhere in Illinois.

All of the project's work in software development, educator training, or information design is geared to supporting specific projects that schools propose as the "price of admission" to the online environment. Projects should be accomplished in a specific time frame, usually three to six months, with clearly delineated milestones agreed upon by all parties. Nearly half of the Collaboratory Project's staff works within the schools. The Collaboratory Project director also presents an annual report to the university and our corporate sponsor, but probably the best evaluation is given by the teachers themselves. Every spring, participants gather for a day-long symposium, where educators can engage in discussion, listen to presentations on particularly successful or interesting projects, and learn from what works and what doesn't. The symposium also provides a great opportunity to start or cement plans for online collaborations. In its first year, the symposium totalled 80 participants. The following year, attendance grew to nearly 200. Anticipated attendance at the 2000 symposium: 300.

At SIGGRAPH 2000, the Collaboratory Project demonstrated some projects in its various communities and reflected upon aspects of its organizational and computing environment that have made success possible, and the challenges associated with further developing the communities. Both the online environment and the project development in the field emphasize the responsibility and commitment of educators engaged in real classroom and curricular projects. This is what has made it possible for the Collaboratory Project to work on a large scale. Some failure is inevitable, but reaching a large audience produces successful projects that serve as models and inspiration for others.

In its third year, the project has reached a critical mass. Its online communities have enough participants that new arrivals can get started by following others' examples. Now the project offers an environment of sufficient depth and scale that it can develop strategic relationships with entire school districts and build an environment where more opportunities for collaboration emerge both in organizational relations and application design. The rapid development of network technologies also offers new opportunities in a wide range of media types. The project is working on streaming audio and video services and planning a more interactive and flexible computing environment through its portal site design. Ultimately, the real challenge lies not in technological innovation but in integration of systems design with the social and material world.



These two images are by students who studied Picasso's art as part of their Spanish class at Murray Language Academy, a Chicago inner city school. They are posting their images and texts to a MediaSpace project they share with Spanish students at Crow Island School in a northshore suburb of Chicago.



techARTS is a community action project that uses the arts, especially in the area of computer graphics, to encourage adolescent girls from the City of Buffalo to learn about and use computers. It began in a pilot phase in the fall of 1999. Development and implementation of the pilot was enlightening, invigorating, and completely surprising. This paper describes development and implementation of both the idea and the actual pilot program, and it summarizes the participants, the evaluation procedure, and project results. It combines the experience of designing and leading techARTS with wider findings from the discussion of girls and technology to argue for an increase in such programs.

The goal of techARTS is to provide all girls, especially girls who have had few opportunities, a chance to create on the computer and interact with role models. The project pays special attention to women using and making technology to provide positive, "technological" female role models to help encourage exploration, creative thinking, problem solving, and originality with a computer.

Numerous scholars note that by the time girls are of junior high school age, they need to know that culture includes women. They need more examples and role models of women who are strong and acting in a variety of roles⁹⁻¹⁰. In addition, research efforts like Rochester Institute of Technology's Project EDGE exemplify the need for more girls and technology programs. Outreach programs can provide a useful introduction to computers^{12,13}.

CONTEXT AND DATA

The multimedia gender gap has attracted some attention (commercial software produced for this age group by Purple Moon, Girl Games, Mattel, and Her Interactive; this work, however, has avoided educational content entirely³). There are no educational technology programs in Buffalo for disadvantaged pre-adolescent girls. Buffalo has both the lowest per-capita income (\$10,445) and the lowest median income (\$18,482) of comparable mid-size US cities (Cincinnati, Tulsa, Sacramento, Saint Louis, Louisville, Albuquerque, Atlanta, Oakland)¹⁵. This project takes place in an economically disadvantaged community. Technology (advertisements, educational CD-ROM products, games) is not often directed at people of color. In order to be successful, this program must infuse technology education with a multicultural approach.

According to a recent US Department of Commerce survey, while the number of African-Americans and Hispanics who own PCs and have access to the Internet is steadily growing, it is still less than the number of whites who have ownership and access⁴. The techARTS

program was designed to appeal to 9-11-year-old girls from a variety of ethnicities and backgrounds. In the planning stage, the project consulted several studies, including the 1994 Community Resource Curriculum Development for Grades 5-6, which emphasizes ways to develop multi-ethnic, multi-cultural curricular units in scientific fields. The authors note that it is important to use skills students have gained from their own lives in hands-on investigations and experiential activities.

THE PILOT TECHARTS PROGRAM PLAN

Logistics

The plan was that approximately 40 girls age 9 - 11 from three Buffalo-area school districts would meet once a week for 12 weeks during the fall of 1999. The goal of the program was to work with computer tools in creation of artwork, which would increase the girls' confidence and competence in using technology. The meetings were held at Buffalo's nonprofit urban media center, Squeaky Wheel/Buffalo Media Resources.

Meetings were scheduled on Saturdays for a practical reason: transportation was donated to the program. The project plan emphasized giving the girls ample time to use the computers, since statistics reveal that girls spend much less time with computers than boys. Times-Mirror and SPPA surveys in 1997 estimated that boys' use of computers at home is 20 percent higher than girls'⁶. To attract girls and extend the time girls spend using the computer, classes were scheduled from 1-4 pm on Saturdays. A full-time computer graphics teacher and two student assistants were recruited to run each week's courses. A nationally known guest woman computer scientist or artist utilizing technology would come to speak once per month, to answer questions about her career and why she chose to work in this field. At the next two sessions, the girls were to do hands-on activities related to the mentors' artistic techniques used or media. The last two sessions featured a final project show and graduation. The girls had an opportunity to try out all kinds of digital techniques, including creating computer graphics, making Web pages, and writing songs in the computer lab.

The Participants

Girls from two low-income neighborhoods surrounding the media center attended the techARTS sessions. The girls were bussed from community centers to the media center. Because the program was entirely voluntary, the sessions had to be attractive to the girls.

The Staff

Studies suggest that to interest girls in technology and foster ethnic identity, program staff should reflect the diversity of the population being served. Many gender-specific programs hire all female staff, and these women are important in providing good role models⁶.

Educational Approach

techARTS was created to deliver models of collaborative, problem-based learning within an active learning paradigm⁷. According to Nelson Ikegulu, true multicultural learning can take place if the curriculum relates experientially to the cognitive, academic, social, and linguistic abilities of students. He argues that learning traits should be considered in computer-mediated program design⁸. Miller further notes that collaborating, rather than competing, offering one-on-one instruction, and providing role models can help get girls interested in technology⁷. Multiple styles of learning were addressed in the program design.

Evaluation/Assessment

The assessment process began in the research stage of the project. Special attention was paid to how the content and curriculum matched the audience. Student perspectives were assessed to determine quality in teaching and learning.

Implementation of the Pilot, Results, and Lessons Learned

What actually happened when the plan was implemented? Almost every aspect of the planned program was altered as the pilot program was implemented:

1. The plan was to gather approximately 40 girls age 9 - 11 from three Buffalo area school districts to meet once a week. In reality, the nonprofit media center Squeaky Wheel possessed three Macintosh computers, and the University at Buffalo grant provided two more PCs. So the program had five computers for 40 girls. Sharing was something the girls had to do at school and at home, and it became clear that some girls would rather sit in the corner than compete for the use of the mouse, even within organized, structured exercises. Eventually, the program served about 12-15 girls each week.
2. While some of the girls had used computers before, none of the girls at the first session could tell the organizers and teachers what computer graphics were. The word "graphic" was unknown to all the participants, but art and drawing images on paper were familiar.
3. Not only girls age 9-11 were interested. Some girls brought older cousins and younger sisters, and guidance counselors or community activists sometimes delivered other students. The program decided not to turn girls away if they did not meet the 9-11 criterion. This did cause problems, however. One 15-year-old girl, a cousin of an 11-year-old, voluntarily tagged along each week but complained loudly during the whole session. In the end, she did enjoy using the computer.
4. Though the plan was to serve girls from a variety of ethnicities and backgrounds, the program actually served primarily African American girls, even though Hispanic and Caucasian girls also lived in the target neighborhoods.
5. Out of approximately 29 different girls who attended the program, only three had access to computers outside of school, and one quarter did not have any access to computers, even at school.
6. The girls had difficulty paying attention for long weekend hours, no matter how enjoyable the class. More importantly, most if not all of the girls came to the sessions without having eaten. Some had no food at home, and others forgot to eat. To run the class with attentive kids, the team had to find adequate lunch food for the participants. Food might be an issue for every program of this type regardless of time of day. The food and drink for 12- 15 participants came to about \$22.00-\$30.00 per week.
7. The program was designed to incorporate collaborative, problem-based learning. Instead, it started with active learning, where the girls were empowered to become independent problem solvers and creative participants. Learning was an active and constructive process.
8. In the first session, the instructor showed the students a computer graphics program. Activities were projected onto a wall with a LCD projector, and the students were asked to assemble materials for a self-portrait. This kind of instruction was useless. Not only did it feel like school, it was Saturday, and participants wanted to talk. When the concept of compositing was introduced, each girl was asked to place her own image inside another picture, then given the tools necessary to create their own ideas. The girls became engrossed. Contextual tasks that related to kids'everyday lives proved to be a great motivator. For example, the image compositing exercise was the weekend after a Lauryn Hill concert, which some of the girls attended. The compositing assignment changed to match this context, so that girls could be on stage with Lauryn Hill or another favorite group in concert. This proved to be better motivation than other exercises or program tasks.

techARTS: Urban Girls Get A Handle on Computer Graphics

9. The mentor part of the program was fantastic. The girls were in awe of the professional women who came to mentor them in computer graphics. One of the mentors, a Latina designer from Austin, Texas, showed the girls how to make logos for themselves. A student employee then showed the girls how to animate their own personal logos. The girls wanted mentors each week! And no matter how talented or renowned, the girls did not see the weekly instructors in the same light as the mentors. Mentoring programs are one way to help girls overcome cultural bias about women and technology.
10. The plan was to have a full-time computer graphics teacher and two student assistants run each week's course. This plan was a good start, but more instructors were needed to keep the girls engaged. Two extra volunteers were recruited to help provide one-on-one instruction. This was useful, but individual instruction is very difficult to manage and afford. Fortunately, the program had dedicated volunteers. Small groups or individuals met with the most success during the pilot. Other kinds of instruction (lecture models to collaborative problem-solving teams without an adult group leader) were not very effective at the introductory level. It was as if one could either turn interest completely on or off. There was little middle ground with this group.

Instructional models changed a bit as the pilot progressed. Once they were trained, one of the more confident girls, could help other girls perform tasks like scanning and drawing and thus form their own small work groups. The first sessions required the heaviest personnel investment, and fewer personnel were needed as students progressed and managed their own groups and tasks.

11. The participants tended to fall into one of four groups during the sessions: girls who wanted to explore the Internet, girls who liked to draw and create graphics, girls who liked to play computer games, and girls who did not want to do anything. But members of the groups changed from session to session. One week, a girl might gravitate towards playing games, but the next she might prefer making graphics. Thus the participants in these four quadrants shifted, but the same categories were observed in almost every session. It was apparent that certain girls were doing different types of activities to meet the needs of different learning styles. The largest number of girls each week were those who wished to visit chat rooms and fan Web sites. This observation reinforces research that suggests that females see technology as a means for social communication.¹

12. Assessment information was to be gathered from the girls through usability tests and a variety of formal and informal surveys throughout the life of the project. But during the middle of a written survey about the course, it became obvious that some girls could not write, some wrote too-short answers, and others purposefully wrote contrary and joking answers to the questions. How could one assess students who did not wish to be assessed, because it reminded them of school? How could one keep girls' interest in technology and art when outside the sessions they face many, many other challenges? One girl's mother had passed away a year before she joined the program, from AIDS-related illnesses. Another was sent from relative to relative and from school to school around the city. It is difficult to imagine the problems until you encounter each individual. So when it comes time for written assessment, things became difficult. Survey questions tended to receive short answers. For example, several students answered "What is the most important thing you are learning, and why?" with "a lot" (no details or explanation). Interviews proved to be more successful. Even with the problems with evaluation procedures, the project did collect much useful data through some of the written evaluations and oral interviews.

Part of assessing the success of the program was the fact that the girls attended. Attendance was free and voluntary. It was clear that parents were typically not involved with the decision to attend techARTS. Only one parent out of approximately 29 different girls who attended appeared at the media center or phoned the program. Therefore, the girl's eagerness to attend was a marker of success. Girls' relative lack of computing experience is particularly significant, as it may diminish their self confidence and performance¹.

Assessment indicated some areas of success, including instructional atmosphere and subject areas explored. The girls liked the informal setting of the nonprofit media center as opposed to a school computer lab. "It feels like my living room, not school," a participant said. They also liked that the program was girls-only, and the instructors and volunteers were all women. "It gave me a chance to do stuff on the computer without fighting with [a] boy," one girl noted. Some responses indicated self-confidence with the computer. One participant responded: "I think boys should be able to come too, so I could talk to them, [and do] less on [the] computer since I know it."

Summary

The directors, mentors, and interested participants planned the pilot by believing that the girls in the techARTS program would have a greater level of comfort and a greater chance of success using technology after attending the program. In the short term, this meant that the group hoped that the girls will continue to attend and create artwork using computers. Those who continue will also be able to mentor others. In the long term, this will result in more women entering high school, college, and careers who are technologically literate, confident, and ready to pursue computer-related interests on the same footing as their male counterparts. The program, with its foregrounding of computer graphics, will also produce girls open to the creative and technical possibilities, girls who would never have had the attention or opportunity to excel in this area.

More programs like techARTS are needed; these programs will clearly be significant for girls around the country, because they offer increased opportunity, and to industry, because they address the national shortage of information technology professionals by attracting and graduating more women fluent in the use and design of high technology and digital media¹⁶.

techARTS is funded by the University at Buffalo and by the RGK Foundation, Austin, Texas.

References

1. Brunner, C. Opening technology to girls: the approach computer-using teachers take may make the difference. *Electronic Learning*, 16 (4), 55.
2. Craig, D.V. & Stewart, J. Mission to Mars: A collaborative project infusing technology and telecommunications into the curriculum. *Learning and Leading with Technology*, 25 (2), 22-27.
3. Fun, but based on products. (1998, May 21). *The New York Times*.
4. Hall, R. (1998, July 29). Blacks, Hispanics still lag behind whites in level of PC ownership, Internet access. *The Wall Street Journal*, p. B12 (E).
5. Ikegulu, N. (1997). Effectiveness of mediated instructional strategies and learning styles in multiculturally linguistic environments: Implications for developmental educators. *ERIC #*: ED409757.
6. Littleton, K. et al. (1993). Gender and software effects in computer-based problem solving. Annual Meeting of the Society for Research in Child Development, New Orleans, LA, 25-28 March 1993.
7. Miller, L. et al. Girls' preferences in software design: Insights from a focus group. *Interpersonal Computing & Technology*, 4 (2). 27-36.
8. *The Office of Justice Programs. What are the elements and features of promising programs?* ojjdp.ncjrs.org/pubs/principles/ch3_3.html October 1998.
9. Orenstein, P. (1994). *Schoolgirls: Young women, self-esteem, and the confidence gap*. Anchor Press.
10. Pipher, M. (1994). *Reviving Ophelia : Saving the selves of adolescent girls*. Putnam Publishing Group.
11. Sackrowitz, M. & Parelius, A. (1996). An unlevel playing field: Women in the introductory computer science courses. *Proceedings of the Twenty-Seventh SIGCSE Technical Symposium on Computer Science Education*, February 1996, 7-41.
12. Teague, G. J., & Clark, V. Attracting women to tertiary computing courses: Two programs directed at secondary level. *Proceedings of the Twenty-Fourth SIGCSE Technical Symposium on Computer Science Education*, February 1993, 208-212.
13. Townsend, G. Viewing video-taped role models improves female attitudes toward computer science. *Proceedings of the Twenty-Seventh SIGCSE Technical Symposium on Computer Science Education*, February 1996, 42-45.
14. *Statistical Abstract of the United States: 1997, 117*. Washington, DC.: US Bureau of the Census.
15. US Bureau of the Census. Online database: www.census.gov/, 1998 .
16. Zehr, M. A. Computer giants look to students: Companies setting up certification programs. *Education Week* 17 (31).

A celebration of youth lies at the heart of the matter. This is reflected in varied explorations of digital media by young people. Their individual and collective inspirations encompass projects that communicate wonder, vision, and perception. The artifacts (digital, ephemeral, and material, as well as memory and relationship) are forms of evidence of social and creative processes surrounding the studio experience of TeleCommunity.

Since 1996, TeleCommunity has been affiliated with the Duquesne University Studio Art Department in Pittsburgh. Prior to that, it was associated with the Duquesne School of Education. Historically, TeleCommunity grew out of early projects at the Pittsburgh Children's Museum and Carnegie Mellon University in 1992. Most recently, TeleCommunity has incorporated into a non-profit organization entitled TeleCommunity: Visible Concept, which is now established as a public foundation. Participants, including teachers, high schools, universities, museums, conferences, companies, consultants, artists, art educators, technologists, and university interns, have all contributed to TeleCommunity projects that involve young people from various cities and countries. These diverse assets and resources have been brought together to form a basis for alternative modes of education, to help realize human potential.

Each year, students 10-18 years of age develop and share their interpretations of themes. Individual and group projects evolve over a period of months, beginning with an original concept, design, or narrative for an animation, computer graphic imagery, an interactive piece, a 3D rendering, or a Web site. Mentors and teachers assist in areas of skill enhancement, conceptual content, and aesthetics. In networked multimedia studio settings, students advance through a range of skills and apply various graphic, animation, interactive, and Web software tools to achieve their goals.

Theme-based work is exchanged over the Internet. New work is undertaken in response to the reactions and interpretations of distant peers. This includes incorporation of shared designs and motifs into digital imagery and recontextualization of narrative material. There is an open-ended aspect to TeleCommunity projects, in that there is freedom to explore and experiment with techniques, approaches, and ideas. Young people work as individuals and together in small groups to realize projects. This team approach offers practice in the principles of collective design, where different participants contribute their skills and abilities, and complement each other.

The project uses a variety of software for 2D and 3D graphics, animation, image processing, and Web production: Photoshop, Painter, Illustrator, Director, Flash, Bryce, Strata Studio Pro, LightWave 3D, Premiere, QuickTime Pro, CU-SeeMe videoconferencing, CyberStudio, Cold Fusion, AOLInstant Messenger, WebCam, RealVideo and QuickTime tools for Internet communication.

The Pittsburgh location uses PowerMacs with high-speed Internet access along with scanners, digital cameras, and color printers. Professional resources, which are beyond the means of many home settings, are placed in the hands of students in advance of the normal school teaching schedule and environment. Internet video-conferencing, digital image exchange, Web-site linkage, and Web broadcasting build friendships, strengthen social bonds, and offer alternative perspectives. The internal energy of social interactions and collaborations on the local scale is shared externally between autonomous groups. Relationships and understanding emerge from this dialogue, creative pursuit, and comparative experience.

The infrastructure goes beyond networks, computers, graphics, and telecommunications software. It includes concepts and practices such as: the power of play to open channels of self-expression; young peoples' willingness to risk growth, to experiment and share inner thoughts and feelings; the value of collaboration in art, design, and education; the catalyst of the creative act in the quest for self-knowledge and identity; and the commitment of a volunteer community of like-minded individuals.

The project's base in Pittsburgh is an outreach effort into the community, with students from local high schools and some home-schooled individuals. Parents play an important role. They bring their children in on weekends and sometimes stay during the studio. All of them are aware of the activities and individual progress of their children. The students step into an informal world full of ideas, expertise, skill, and expressive energy they can draw from and contribute to. Projects have focus. Individuals follow their strengths and interests, and they give advice on direction and pursuits for TeleCommunity as a whole. Select students are responsible for keeping everyone informed about TeleCommunity by maintaining our mail server and updating our Web site.

The staff includes teachers, artists, interns, and consultants in computer graphics, 3D modeling, multimedia, Web design, and museum-based youth art education. Projects have included exchanges of computer graphics, ideas and experiences among students in Pittsburgh, New Orleans, Los Angeles, Jerusalem, Istanbul, San Francisco, Orlando, Norway, and Japan. Our Web site (www.telecommunity.org) includes student work, background material, links to partners, and an area for online file exchanges and online messages. Overarching themes such as: "Illuminations", "Where Stories Meet", "Circle of Time", "Heroes", and "Journeys" have provided conceptual frameworks for ensuing studio efforts and exchanges of digital work.

Partners have different teaching environments, and students are involved in different programs at each site. There are parallel interests and a flexibility within each setting for participation on TeleCommunity projects. Activities include Internet videoconferencing, image and animation exchange, email, and Web-site linkage.

Examples of collaborative settings past and present:

- In the New Orleans area, there is a classroom situation throughout the year at Metairie Park Country Day School. Artist and teacher Carrie Lee Schwartz directs efforts there.
www.telecommunity.org/Horizons/new/rightpgnew.html
www.telecommunity.org/Stories/stories-web/NO/tc-no.html
www.telecommunity.org/circle/neworleans
- Students in Jerusalem take studio classes at the Israel Museum Youth Wing as part of the Multimedia Education Unit program. Susan Hazan, Curator of Multimedia Education and other teachers oversee courses there.
www.telecommunity.org/Horizons/jer/rightpgjer.html
www.telecommunity.org/Circle/Jerusalem/jer99.html
www.telecommunity.org/Jerusalem_Art/JerusalemArt.html

In addition, a small group of young Jewish, Muslim, and Christian students work together on related art projects at the museum. They have made a connection to TeleCommunity in the past.
www.telecommunity.org/Jerusalem_Peace/JerusalemPeace.html
www.telecommunity.org/Stories/stories-web/EastJerusalem/paley.html
- In Istanbul, there are classes at the Bilfen Private Primary School. Children who are learning English write stories and create graphics with teacher Gulsev Sertoglu.
www.telecommunity.org/Circle/Istanbul/ist99.html
www.telecommunity.org/Stories/stories-web/Istanbul/ist-image.html

- In Pittsburgh, the home-base, middle school and high school students are invited to visit Duquesne University on weekends. They work out of the university multimedia labs. Robert Dunn, Melanie Carr, Sim Graves and Duquesne University student interns teach there.
www.telecommunity.org
www.telecommunity.org/photo-portraits/index.html
- At SIGGRAPH 95, SIGGRAPH 96, SIGGRAPH 97, and SIGGRAPH 98, local students visited our onsite networked studio exhibits. For one week each year, students were immersed in hands-on activities, developing computer graphics and animations, enhancing their skills, and sharing with remote partners.
www.telecommunity.org/Stories/stories-web/Orlando/orl.html
www.telecommunity.org/Horizons/la/rightpgla.html
www.telecommunity.org/Journeys/NOWEB/NewOrleans.html
www.telecommunity.org/electric_garden/index.html
- In Japan, Keigo Yamamoto is director of World Wide Network Art project which has linked to TeleCommunity during the SIGGRAPH conference.
www.syba.co.jp/k-bit/
www.telecommunity.org/Journeys/Kbit.html

Contents

- | | | | |
|----|---|----|--|
| 22 | Achieving Balance: The Challenges of Developing Web-Based Courseware | 51 | Painting Together Across the World: Collaborative NetPainting Between High School & Professional Artists |
| 24 | Beyond The Walls: Redefining the Museum as an Educational New Media Publisher | 54 | A Project-Oriented Graphics Course |
| 26 | A Collaborative and Interdisciplinary Computer Animation Course | 56 | Projects for a Computer Graphics Programming Course |
| 30 | A Crash Course in Creative Digital Media for the Community | 60 | QuickWorlds: Teacher-Driven VR Worlds in an Elementary School Curriculum |
| 34 | Developing a Curriculum on the Arts and Mechanics of New Media for an International Post Graduate Audience | 63 | Teaching Computer Graphics in Brazil: Social Commitment, Creativity and Passion—Against All Odds! |
| 40 | Enhancing Teaching and Learning with Animation | 66 | Teaching Computer Graphics with OpenGL |
| 41 | Interaction in an IVR Museum of Color | 67 | Using 3D Animation in Teaching Introductory Electromagnetism |
| 45 | Is There Any There There? Design and Implementation of Successful Team-Oriented Educational Environments for Computer Animators | | |
| 47 | A Multidisciplinary Approach to Masters Education in Digital Research | | |

E D U C A T O R S P R O G R A M

VALERIE A. MILLER
GEORGIA STATE UNIVERSITY



JURY MEMBERS

Colleen Case
SIGGRAPH 2001 Educators Program Chair
Schoolcraft College

Harry Dangel
Georgia State University

Valerie A. Miller
SIGGRAPH 2000 Educators Program Chair
Georgia State University

Jacquelyn Ford Morie
USC Institute for Creative Technologies

Adele Newton
SIGGRAPH 2000 Community Outreach
Chair
University of Toronto

REVIEWERS FOR THE EDUCATORS PROGRAM

Doug Acheson

Eglin Ayson

Marc Barr

Debi Barrett-Hayes

Robert Blystone

Glenn Campbell

Colleen Case

Steve Cunningham

Harry Dangel

Gitta Domik

Dena Elisabeth Eber

Jacob Furst

Frank Gladstone

Deborah Howes

Joe Lohmar

Francis Marchese

Valerie A. Miller

James Mohler

Barbara Mones

Jacquelyn Morie

Adele Newton

Scott Owen

Jonah Peretti

Maria Roussou

Steve Schain

Jill Smolin

Chris Stapleton

Karen Sullivan

Scott Wilson

Rosalee Wolfe

ADMINISTRATIVE ASSISTANT

Tom Burkhart
University of Iowa



"The principle goal of education is to create men who are capable of doing new things, not simply of repeating what other generations have done —men who are creative, inventive, and discoverers."

JEAN PIAGET

In August 1999, the SIGGRAPH 2000 Educators Program issued a call for submissions that embraced the basic tenets of not only teaching computer graphics and using computer graphics to facilitate learning across the curriculum, but also innovation, creativity, diversity, and inspiration. To say we were successful in our solicitation would be an understatement. SIGGRAPH 2000's Educators Program is filled with all of these things and more.

Through collaboration and cooperation, our contributors have envisioned and developed pedagogically sound courses to provide students with rich, interdisciplinary scenarios for learning. These visionary presentations range from:

- Utilizing virtual reality in the elementary classroom to exploring the curricular, technical, and interactive design challenges in developing Web-based courseware for secondary education.
- The results of collaboration and partnership between the Digital Design Studio at the Glasgow School of Art and businesses, schools, and community to developing a curriculum on "The Arts and Mechanics of New Media" for an international post-graduate audience.
- Projects for a computer graphics programming course that serves effectively as both a service course for the sciences and an introductory course for computer science to the pedagogical approach, curriculum, and results of developing an interdisciplinary course in computer animation for artists and computer scientists.

Each of our submissions was reviewed by highly qualified individuals from academia and industry. After the review, each submission was considered by the SIGGRAPH 2000 Educators Program Jury, a panel of artists and scientists, academics and industry persons. What we present here is but a snapshot of computer graphics in education today and an extraordinary vision of our future.

"Education is a social process. Education is growth. Education is not a preparation for life; education is life itself."

JOHN DEWEY

INTRODUCTION

The mission of SAS inSchool is to provide educators and students with high-quality, content-rich, interactive, Web-based courseware that effectively conveys topics and concepts. We know that computers can enhance classroom instruction and learning in each of the major curricular areas, if the programs are well-designed from both a pedagogical-curricular perspective as well as a technical design view. Our focus is strictly on education, not “edutainment.”

PANEL DISCUSSION TOPICS

Translating this mission into reality presents numerous challenges:

- Curricular goals are sometimes at odds with, or can not be met by, available technologies.
- Other constraints require team members to compromise on the original product vision.
- Assumptions about implementation and interface design.
- The collaboration and communication that is required throughout the development process.

A curriculum specialist discusses the challenges she faces, including how to identify content that is suited to the medium, as well as content that does not merit multimedia treatment; how to translate traditional classroom instruction into a more effective multimedia experience; and how to revise treatment and/or scope without compromising either the content or the product’s educational value.

A software developer examines the technical considerations that had to be addressed before product development began, as well as those issues that emerged during production cycles. Decisions regarding computing platforms, appropriate use of available technologies, and reliance upon plug-ins were made early on, and had implications for later courseware design. As products were developed, the advances in, and idiosyncrasies associated with, various technologies resulted in the inevitable challenges of integrating product components, responding to the quirks of new third-party product releases, and maximizing performance.

Finally, an interactive multimedia designer reviews the design considerations for each product, including identification of the target audience, establishment of the artistic vision, design of the interface, and use of appropriate and effective multimedia. It is important to recognize that the overall design of a product determines its effectiveness in communicating with the user. The visuals that a designer creates help set the tone for a project and enable the user to establish a connection with the content. Learner characteristics, such as special needs, intellectual ability, and learning style, place valid constraints upon design that must be addressed without negatively impacting the product’s aesthetic value.

The panel addresses specific curricular, technical, and design issues, including:

- How should or can student assessment be handled within a unit of study?
- How can courseware be designed to match multiple potential uses?
- Should answers be provided for student exercises, and should they be password-protected?
- How can quality video and audio be delivered via the Internet?
- Can voice recognition technology be integrated with the Internet?
- How can a large amount of content be navigated most effectively?
- How can interface design optimize the readability of onscreen elements?

Deborah L. Gray

Deborah Gray has nine years of classroom experience. She has taught chemistry, physics, and numerous science electives, and she was science department chair at Enloe High School in Raleigh, North Carolina, a magnet school for academically gifted students. She earned her BS and MAT degrees from the University of North Carolina at Chapel Hill.

Ms. Gray is a curriculum development specialist at SAS inSchool, a division of SAS Institute Inc. Initially, she developed the science and health curricula for Cary Academy, a private secondary school. Currently, she serves as a science content specialist for the division’s production of multimedia courseware for secondary education.

Using her content knowledge and understanding of educational principles, Ms. Gray provides leadership in identifying those science topics that are the best candidates for digital technology support. She has defined manageable templates and developed content treatments that make effective use of multimedia.

“As a high school chemistry teacher, except for modest attempts to develop interdisciplinary units, I had sole responsibility for the instruction that my students received. Determination of content, how it was to be presented, what the student’s role in the learning process needed to be, and how student achievement was to be assessed, were all decisions that I was responsible for making.

“I have been faced with many challenges in making the transition from the classroom to the educational software arena. Successful product development requires a different approach to ‘teaching,’ and my new responsibilities often necessitate my assuming a student role, in order to develop new skills and understanding.

“In my role as a curriculum development specialist, working on Web-based courseware for the secondary education market, I have had to relinquish much of the control, which I used to have over the educational process, to my colleagues. In so doing, I have been afforded the opportunity to work in a collaborative environment where the end product reflects the synergy that results from the melding of individual talents and perspectives.”

Jim Begley

Jim Begley is a software developer with the SAS inSchool division of SAS Institute Inc., where he has focused on science and math products, system architecture, and technology evaluation. He has also developed visualization software and graphics interfaces at SAS. He previously worked at Data General Corporation, where he developed graphics toolkits and computer-aided design tools. He earned a bachelor's degree in information and computer science from the Georgia Institute of Technology in 1983 and a master's degree in computer science from the University of North Carolina at Chapel Hill in 1993.

“During my career as a software developer, I've had the pleasure of working on interesting projects in a number of areas: computer-aided design, molecular modeling, 2D and 3D graphics APIs, and data visualization. Web-based educational software, however, offers a unique set of technical challenges.

“First and foremost, we must realize that technology is a tool for conveying the educational content. Content is king, and all the technologies chosen for use in a product must serve the king. The coolest and latest technologies have no place in our products if they do not help us achieve our educational goals.

“For software to be used in secondary education, we must recognize that the intended market's computing platform is as diverse as its level of computing expertise. We must be careful that our software runs in a variety of environments and is easy to install, maintain, and update.

“Many of the Web technologies are new, so we accept the challenge of gaining experience with the tools while simultaneously trying to create production-quality software. And since these Web technologies change rapidly (many are made available before they are truly sound), we must be flexible and ready to evaluate the latest releases for bug fixes, new features, and stability.

“Many of the technologies are very good at doing one thing, like displaying text or playing video or animating graphical elements, but they must be made to work together as part of a cohesive educational package. We must get disparate technologies from different (and sometimes competing) companies to work together well enough to create the impression of a unified whole.

“Finally, we've got to be sure that our products perform at the level our customers expect at their installation. Tuning and optimization are important steps in the development process that can be easily be put off until it's too late.

“By meeting these technical challenges, we can create unique educational products that illustrate and illuminate.”

Deanna Scinta

Deanna Scinta earned a master's degree in computer graphics design and an advanced certificate in interactive media design from the Rochester Institute of Technology. She has more than six years of experience developing interactive CD-ROM titles, corporate Web sites, and kiosk presentations. One of the products she co-designed was CatWorks, an educational CD-ROM of a digital cat dissection for secondary school and college students. Since 1997, she has worked as an interactive media designer at SAS inSchool, creating Web-based courseware for secondary education. Currently, she is the lead designer for English courseware, the company's external Web site, and school administration applications. She also played an integral role in the development of the SAS inSchool corporate identity.

“As an interactive media designer on product development teams at SAS inSchool, I create Web-based courseware for secondary education. Through this experience, I have encountered the design challenge of creating an artistic vision for content developed by curriculum specialists. This challenge was complicated by the complex needs of the target audience and their often-limited access to high-end technology. We addressed this problem through interface design, by utilizing appropriate multimedia techniques within the technological constraints.

“To resolve these issues, it was necessary to have a strong understanding of a product's content. This helped to achieve the aesthetic tone of the product, to develop a relationship between the content and interface design, and to determine appropriate usage of multimedia to not only engage users, but to also attain the desired learning objectives.

“In addition, it was essential to maintain open communication among team members during development cycles. At times, compromises were made to refine the product vision, meet deadlines, and adapt to technology constraints.

“As a result, the development team has created dynamic interactive products that continue to evolve into a unique final form. As an interactive media designer, it is my job to rise to new challenges as we enrich our product line.”

Beyond the Walls: Redefining the Museum as an Educational New Media Publisher

Panelists
Gretchen Schwarz
American Museum
of Natural History
frauline@amnh.org

Patty Abt
American Museum
of Natural History

As museums continue to adopt new media technologies, their role in national education has expanded from traditional research, collection, and exhibition to include original educational publishing using every flavor of new technology. These new publishing initiatives extend the museum's influence beyond its walls and generate fresh production challenges. As publishers, museums have a tremendous opportunity to make a valuable contribution to the available media at a time when raising the standard of science and mathematics education is a national imperative. Though many museums are not seasoned new-media publishers, they do have a natural wealth of content resources waiting to be sculpted. How will museums become important educational publishers? How will they use the most current technologies to deliver their products to reach a significant audience? What technologies are most suitable for crafting learning tools?

In 1997, the American Museum of Natural History in New York began to tackle these questions with the launch of its National Center for Science Literacy, Education, and Technology. This panel presents specific examples of the National Center's current developments in new media education publishing: the Ology Web site and Science Bulletins. Panelists outline the unique technology base that supports each project and discuss practical production challenges. The panel demonstrates that the most effective uses of technology in development of editorial content are not necessarily in costly innovation and infrastructure, but rather in reinterpretation of traditional museum roles in education and the community that support long-term institutional goals. Panelists address how each publishing project exemplifies one of the following production goals:

- Promote experiential, life-long learning (the Ology Web site: www.ology.amnh.org).
- Use dynamic current events as a focal point for introducing and reinforcing major concepts in science curriculum and curatorial content: Science Bulletins

PATTY ABT TOURS THE OLOGY WEB SITE

A virtual museum visit can never replicate the awe a child feels the first time she stands in front of a T-Rex skeleton or wanders through a hall and stumbles upon a model of a Sioux Indian teepee. On the other hand, a Web site can invite visitors into places they could never go before and introduce them to the scientists who explore there. For 130 years, the American Museum of Natural History has been a place of wonder for kids, where, around every corner, they can bump into something mysterious that inspires them to look at their world in a new way. Ology, the Museum's new Web site for kids, is created to extend this relationship beyond the walls of the physical museum to kids across the country.

The underlying premise for Ology is that kids are naturally motivated by a desire to solve mysteries and acquire knowledge. Using a conventional Web site format with additional features like "virtual collecting," the Ology site capitalizes on this innate desire to engage kids in science. The Ology Web site's frequently updated interviews, feature stories, and games are organized around -ologies such as arachnology or paleontology, topics that kids love to dig into. Ology is responsive to kids' interests and desire to participate. It includes clubs to join, polls, interaction with scientists, and a place to help build future -ologies for the Web site.

As kids explore Ology, they collect Ology cards about ideas, people, places, or events that they want to put aside in their own personal area to remember. This innovative technology also allows kids to organize and display cards on the Ology site. Instead of simply sifting through content, as kids do on most Web sites, the Ology cards put kids in the driver's seat, allowing them to choose what topics to explore, create original interpretations, and share them with others. Many kids feel that science is out of their reach, abstract, scary, or boring. The Ology cards give kids an opportunity to be successful in science, to feel that they have mastery over the content because they can interact with it.

GRETCHEN SCHWARZ OUTLINES
THE SCIENCE BULLETINS INITIATIVE

There is no question that the Internet and new technology have broadened access to current events information to 24 hours a day and seven days a week. Museums that collaborate with educators and curriculum to develop educational tools now must add the immediacy of current events. More and more, museum digital tools are being seen on television, the Web, DVD, and interactive productions alongside news, and entertainment programming. There is enormous opportunity for museums to grab visitor imagination by bringing regularly updated news into the exhibition space and to explain "the rest of the story" (the science and history behind the news) using dynamic combinations of images, animations, text, and research links. Interpretations that make vital connections to curatorial goals and curriculum form the Museum's unique editorial voice.

The Science Bulletins production team at the American Museum of Natural History began integrating current events into exhibition halls in 1999. Science Bulletins produces interactively linked articles for touch-screen kiosks, high-definition television graphics, animations, short video documentaries, and linked Web sites for installations in three of the new permanent exhibition halls. The subjects are biodiversity, earth science, and astrophysical science.

New technologies make presentation of frequently updated current events material sustainable with a relatively small production staff. A team of eight, plus consulting scientists and researchers, steer the editorial and technical production of daily programming. Resources and production process include:

- Partnerships with news organizations that provide video feeds over fiber-optic lines.
- Graphic animation templates created with Everest, a real-time render authoring tool.
- Networked touch-screen kiosks with customizable template structures.

In 1997, The American Museum of Natural History, in an initiative with NASA, launched its National Center for Science, Literacy, Education, and Technology to create original science educational publishing for expanded audiences. Using the World Wide Web, interactive television, books, classroom curricula, and more, The National Center creates programs and materials that connect people of all ages nationwide to the unparalleled resources of the museum collections of 32 million specimens and artifacts, 43 exhibition halls, more than 200 research scientists, and 128 years of expertise in educational programming. In 1998, the National Center added the Science Bulletins team to provide dynamic current-events science reporting of "the science behind the news."



Breaking news on the Science Bulletins Astro Bulletin wall.

A Collaborative and Interdisciplinary Computer Animation Course

David S. Ebert
University of Maryland
Baltimore County
ebert@umbc.edu

Dan Bailey
University of Maryland
Baltimore County

This paper presents the pedagogical approach, curriculum, and results we used to develop an interdisciplinary course in computer animation for artists and computer scientists.

INTRODUCTION

Computer animation has always required a close collaboration between artists and computer scientists. However, traditional educational approaches to computer animation isolate artists and computer scientists from each other. Traditional courses fail to teach students the important interdisciplinary nature of computer animation and fail to provide practical interdisciplinary collaborative work experience to students. We have developed an interdisciplinary computer animation course that focuses on contemporary issues in computer animation and requires the skills of animators and programmers working in teams. The goals of this course are to:

- Develop students' teamwork and effective group dynamic skills.
- Increase the technical graphics and animation knowledge of computer science students.
- Increase the animation skills and knowledge of advanced computer animation techniques of art students.
- Introduce art students to the technical aspects of rendering and animation, and expose them to research issues in computer animation.
- Introduce computer science students to traditional and computer animation techniques.
- Introduce art students to the creative potential of writing procedural shaders, models, and animation expressions.
- Provide practical animation production experience, using and extending commercial animation software.
- Provide a collaborative learning environment where students will learn from each other, as well as the course instructors.

A key aspect of this course is that students gain experience in participating in interdisciplinary teams. Teams of four-five visual arts and computer science students work together to produce animations that utilize each member's skills and interests, in a manner similar to commercial animation environments. The computer animation industry requires employees to work in teams on large projects, where the teams are composed of members from quite disjoint backgrounds. We have structured this course to help students learn how to communicate, work, and even thrive in this environment.

BACKGROUND AND MOTIVATION

Animation's history, from its origins in the 1880s to contemporary time, is a continuous line of technological inventions that have allowed animators the ability to achieve higher-quality effects with greater ease.¹ At the heart of each of these developments has been the successful synthesis of artistic and scientific talents, usually

resulting from creative collaborations. Equally as important, large-scale animation production has always required large teams of variously talented individuals. The Warner Bros. animators of the 1930s through the 1950s produced some of the most successful cartoons of this century, resulting from effective collaboration among its directors, animators, writers, technicians, artists, and musicians.¹

Both of these issues point to the fact that a successful and contemporary animation curriculum should not only be interdisciplinary, but it should also encourage students to develop effective team skills. However, crossing the boundaries between areas and departments within universities has always been difficult. Therefore, many schools are slow to address the industry trend to teach and encourage effective teamwork and collaboration between animators and computer scientists. Currently, many animation, special effects, and computer graphics houses are creating their own in-house workshops and programs to address these issues.

PEDAGOGICAL APPROACH

Our pedagogical approach to this course has two key themes: interdisciplinary work and collaborative education. This is true even in the instruction and design of the course, which is team taught by a visual arts faculty member and a computer science faculty member. Most of our lectures are designed to have sections that both faculty members present, highlighting the technical computer graphics aspects and the art and animation aspects of the material. In every aspect of the course, we encourage students to collaborate and help each other. Initial assignments, described below, are designed to provide easy and difficult segments for both types of students, thus encouraging students to begin interacting with their counterparts.

CURRICULUM

The curriculum for the course was developed with three concerns in mind:

1. How to introduce mathematical expressions and scripting to the art and animation students and, at the same time, introduce the computer science students to basic animation concepts and the animation software. This was accomplished through individual projects and tutorials that students completed within the first two weeks of class.
2. How to present the basic aspects of effective collaboration among individuals and create an environment that encourages strong teamwork. This was accomplished through presentations, discussions, examples, and two initial assignments that "broke the ice" and encouraged the teams to utilize the skills of individual team members.
3. How to present both the computer science and art concepts of 3D computer animation. The course was always taught collaboratively by the two instructors. Topics were presented as lectures and demonstrations with each instructor presenting material.

We structured the course to include the following material:

- Group dynamics and collaborative teams.
- Basic mathematics for computer graphics and animation.
- Computer science basics of rendering, shading, anti-aliasing, and procedural techniques.
- Photorealistic rendering techniques and tricks using a commercial animation package.
- Scene composition, camera angles, basics of set lighting, and cinematography.
- Animation basics: timing, keyframing, hierarchical animation, interpolation, and motion paths.
- Kinematics and inverse kinematics.
- Object deformation, soft objects, and blending.
- Dynamics, physics-based animation, and inverse dynamics.
- Behavioral animation: particle systems, crowds, flocks, and emergent behavior.
- Procedural animation and intelligent characters.
- Animating natural phenomena.

We have chosen to use a text for the class³ that is intended to teach computer animation to artists and recommend several computer-science-oriented animation books as supplemental material.⁴

GROUP DYNAMICS AND COLLABORATIVE WORK

Traditionally, students are not given the skills to work effectively in collaborative teams. To address this, we present the basic concepts of effective collaborative work and emphasize this concept by leading a class discussion that relates individual experiences in group work. Students are given background material on teamwork, “groups versus true teams,” and a worksheet to help evaluate the effectiveness of each of their team meetings and identify problems as early as possible. Another important aspect of effective student teamwork is close inter-

action of the teams and the instructors, which allows instructors to serve as observers, help identify team problems early, and then work with teams to correct these problems. Our experience shows that meeting every 10-14 days with individual teams is effective in improving the performance of the teams. Another technique that aids the creation of effective teams is to allow the teams to change after an initial assignment, providing students with more control over the team composition. In our experience, we have not had any teams change during this “free-agency” period, but this opportunity has encouraged teams to evaluate their effectiveness and discuss and resolve problems early in the class.

STRUCTURE AND IMPLEMENTATION

The 15-week course utilizes the Maya software package from Alias|Wavefront. The students start working in teams of four to five students during the second week of the class. Students may propose the composition of their team, but each team must consist of at least two art and two computer science students. The students have an icebreaker team project to perform in Weeks 3-5. We also have some initial assignments that expose artists to working with vectors, angles, and simple illumination, expose computer scientists to key-frame animation, and both to the procedural, extensible aspects of the Maya modeling, animation, and rendering package.

To expose students to the power of procedural shading techniques,² the flexibility of the Maya Embedded Language (MEL) scripting language, and the challenge of photorealistic image generation, the students’ second major project is to generate a photorealistic image or animation that must contain a specific type of light and a specific object element chosen randomly. We also require students to create a GUI for the artists to use in creating the final image or animation. Students are given three weeks to complete this assignment.

The main component of the course is a 12-week team project to produce an interesting, professional animation that includes extending the Maya package to implement a new animation capability. Students must also consider the subject matter of the project and choose an




Figure 1: (a) Sample bluescreen image.

(b) Example final animation still by CSmART Allstars.

appropriate structure and format, since a successful animation is only as good as its story, premise, or content. Early in the semester, each team is required to “pitch” its animation in a professional presentation with storyboards, charts, and slides. This project also requires a presentation of progress to the class after five weeks of work and a final presentation of results to the class. These are graded on presentation, artistic merit, and technical merit.

INITIAL RESULTS

We have successfully offered this course twice. The students enjoyed the class and gained valuable experience in computer animation and working in teams. Most of the teams worked successfully together, and team dynamics varied widely. In fact, one of the most successful teams started out with very poor group dynamics. By the end of the semester, they learned how to communicate effectively, appreciate each other’s skills, and work as a team. Only one team suffered significantly from group dynamics (one team member dropped the class).

Blue-Screen Project

Our ice-breaker team project required each team to composite live-action video with computer-generated effects for a 10-second animation. Teams were required to generate both CG foreground and background elements. This assignment proved to be successful in providing the teams with an initial low-risk and fun project that encouraged creative collaboration. An example still of the blue-screen video is shown in Figure 1a. An example still from a completed project is shown in Figure 1b.

Script and Rendering Assignment

This assignment was presented in two parts. The first part emphasized the computer science students’ skills by requiring each team to create a MEL script and GUI that allowed an animator to create pizzas with different types of crust and different amounts of toppings. An example GUI can be seen in Figure 2a. The second part of the assignment emphasized the artists’ skills in creating a photorealistic image and/or animation of the pizza in a given setting with a specific type of light

source. These settings and light sources were chosen randomly and included romantic dinner by candle light and late night TV dinner by television set. An image showing a photorealistic rendering of a pizza illuminated by a television can be seen in Figure 2b.

Semester Animation Projects

The semester animation projects included a wide variety of technical and artistic styles. Below are summaries of a few team projects:

- *The Autonomous Chicken Farm* by The A-Team. This team explored behavioral animation by developing procedurally animated “creatures.” The example creatures (chickens) had controllable personality attributes (hunger, beauty, and aggression) based on the seven deadly sins. The final product of this project was a MEL script for creating autonomous “smart” creatures and controlling their movement, generating input for the inverse-kinematic controls and expressions that the animators used for controlling the motion of the creatures. A short demonstration of the Maya extension was also produced. Results can be seen in Figure 3a.
- *Expressive Avatar: An Expression-Driven Emotive Artificial Intelligence* by Digital Macabre. This project created a plug-in that uses facial morph targets and an IK skeleton to enable a character to react to a random environment based on programmable personality. The team also produced a computer-animated short film that demonstrated the expressive facial animation of their character reacting to typical events in a hospital waiting room (Figure 3b). The character’s head was modeled from scanned photographs of a student, which were also used to generate the texture maps for the head.
- *Midnight Music* by the Primary Elephants. This project created a plug-in that enables artists to use music to drive their characters’ animations. For the demonstration of this plug-in, several toys were animated to different musical segments. The animation of each toy automatically adapted to different musical segments based on the dance style created for their toy type.




Figure 2: (a) Example pizza creation GUI.

(b) Example final rendered pizza illuminated with TV light by the Digital Macabre.

Many teams completed their projects, but a few did not. Two teams continued their projects the following semester by working independently with the instructors.

Evaluation of Collaborative Teams

To evaluate the effectiveness of the teams, each student filled out a questionnaire at the end of the class. Students were asked to evaluate the amount of work they contributed to each project, which aspects of the team worked well, which aspects of the team didn't work well, and what would have helped the team work better. Comparing these evaluation forms showed that the teams worked very well together. Even the team with the worst performance agreed as to the cause of their team's weak performance. All team members realized that more frequent team meetings and better communication among team members would have made their team stronger.

CONCLUSIONS

We have developed a successful interdisciplinary course to teach computer animation to computer scientists and artists based on interdisciplinary collaborative work. This approach for education is very powerful and rewarding, but it does require a significant amount of effort in teaching not only computer animation but also successful teamwork and group-dynamic techniques. From our experience, the most important aspect of teaching a team-based course is helping students work effectively in teams. Interactive class discussions, frequent meetings with individual teams, and effective role models are very beneficial.

We believe that our collaborative teaching approach provided a good example to the class of the strength and effectiveness of interdisciplinary collaborative work. Our experience has also shown that both the artistic and computer science aspects of computer animation can be effectively taught in a semester course. Finally, our experience has shown that this collaborative, interdisciplinary approach to teaching

computer animation provides more benefit to students than a traditional computer animation course. Students not only learn from the instructors, they also learn even more from their team members, especially from those who are on the "other side" of computer animation.

We would like to thank the student teams who helped us develop and improve this course:

The A-Team: Tracy Corder, Will Gee, Mike Keeseey, Joe Romano

Analog Blacksmiths: Chris Esposito, Mary Levy, Phadke Pratik, Evan Williamson, Damion Wilson

Beasts: Kim Harrington, Mike Madison, Chris Morris, Brian Resurreccion, Shawn Wood

CSmART Allstars: Ava Collins, Alex Eller, Jason Lubawski, Marlin Rowley, Christian Valiente

Day 8: Drew Cobleigh, Kris Kuhn, J McBride, Divyesh Shukla

Digital Macabre: Jeremy Dobrzanski, Dennis Moellers, Dov Horowitz, Mike Sharp, Stephan Sherman

Primary Elephants: Andy Cedilnik, Jodi Kravetzker, Sushama Prasad, Aaron Weidele

Screaming Nixons: Eun Baek, Jon Feibelman, Costas Kleopa, Vlad Korolev, Steve Matuszek

Wookie Pimps: Michelle Hunt, Steve Jacobs, Chris Slingluff, Joy Saunders

References

1. Maltin, L. (1980). *Of mice and men*. New York: McGraw-Hill.
2. Ebert, D., Musgrave, F.K., Peachey, D., Perlin, P., Worley, S. (1998). *Texturing and modeling: A procedural approach* (2nd ed.). AP Professional.
3. O'Rourke, M. (1998) *Principles of three-dimensional computer animation* (Rev. ed.). New York: W.W. Norton and Company.
4. Watt, A & Watt, M. (1992) *Advanced animation and rendering techniques, theory and practice*. New York: Addison-Wesley.



Figure 3: (a) A still from "Expressive Avatar" by Digital Macabre showing responsive facial animation to a character's environment.

(b) A still from the "Autonomous Chicken Farm" by the A-Team showing behavioral procedural animation of articulated skeletons.

“Man is made for something better than disturbing dirt, and all work of that kind should be done by a machine.”

OSCAR WILDE

INTRODUCTION

As we race into the third millennium, we can look back on the 20th century as being revolutionary in many areas, not least in the development of new technology designed to improve the human experience. The emergence and rapid development of computing technology has led to automation of many mundane repetitive tasks, leaving more time for creative thought and leisure activity on the one hand and increases in productivity on the other.

The age of information is in full flow. In Scotland, the economy has moved from heavy industry, such as shipbuilding and engineering in the first half of the last century, to information technology and electronics in the latter. Silicon Glen has taken a firm grip on the knowledge-based economy.

In Glasgow, one of the longest-established creative institutions is the Glasgow School of Art. Opened originally as the Government School of Design in 1845, the school was re-housed in 1899 in its current home, the Charles Rennie Mackintosh-designed Glasgow School of Art building. The school has continued to contribute to the growth and development of the city throughout the 20th century.

THE DIGITAL DESIGN STUDIO

The most recent development at Glasgow School of Art is the establishment of the Digital Design Studio (DDS). Situated in The House for an Art Lover in parkland on the south side of Glasgow, the DDS is the Glasgow School of Art postgraduate research centre. The Studio is equipped with a range of SGI workstations from O2 to Onyx2 Infinite Reality running high-end 3D animation software from Alias|Wavefront. The studio is the hub of new research into design-centred creative digital technologies and applications in partnership with large corporations including the Ford Motor Company, the BBC, British Telecom, and Scottish Enterprise.

In addition, Glasgow School of Art is the lead institution in a collaborative research venture involving Glasgow University, Strathclyde University, and Caledonian University. The Virtual Design Institute combines the diversity of expertise in these organisations to investigate and develop rapid and virtual prototyping.

The DDS also houses the two-year multidisciplinary masters course in 2D/3D motion graphics and virtual prototyping, involving students from a wide range of backgrounds, from psychology, marine engineering, and microbiology to product design and fine art. The requirement for entry stresses creative ideas to advance the field rather than expertise in computing. The students are taken through an intensively taught first year: a crash course in creative computing. From this platform, they move into their second year with the basis of an advanced skill set to tackle their personal research projects.

FEEDBACK

Establishment of advanced research centres in academic institutions is hardly a new concept. With cutting-edge technology in a high-tech environment, it is not surprising that both leaders in industry and government are quick to support an initiative such as the Digital Design Studio. We are almost guaranteed a steady flow of talented masters students to take up sought-after places to stretch and develop their ideas with the new digital technologies. The opportunities for these new pioneers are almost limitless. However, if the next century is to build on the achievements of the last, we must cast the net wider and look towards providing access to sections of society previously excluded from the immediate benefits of harnessing creative digital technology.

DIGITAL EVANGELISM

Many key developments in digital design take place behind closed (often locked) doors, where information is available on a strict “need-to-know” basis, protected by non-disclosure agreements. The results of this activity only enter the public domain when the commercial benefits have been safely tied up, often years later.

There is a distinct danger of a technological divide developing within our society where the digital “haves” actively participate in future developments and the “have nots” become even more disenfranchised.

How then do we encourage a new generation to become involved in the digital revolution when the public image is often out of date, and the means to become involved are limited by access to the technology itself? We must start to feed back the experience gained from our research into the rest of the community at the earliest opportunity. The first step in this process is through education.

While the DDS is committed to working in a number of demanding and highly sophisticated research projects as well as the education of masters students, spreading the word about digital technology to other groups in society is of great interest. It is all too easy for individuals to be dissuaded from involvement in advanced technology for reasons of language, expense, and that unfortunate “users club” into which digital artists and designers can easily segregate themselves.

In a series of initiatives aimed at breaking down these barriers, links have been established with local community groups with the general aim of revelling in digital imagery. As good fortune would have it, the studio was funded to produce an explanatory animation aimed at the young. This evolved into “The Tail that Wagged the Dog,” an interactive attempt at explaining in simple language what is involved in producing a digital animation. Similarly, the opportunity to work with a budding 12-year-old automotive designer was taken to demonstrate how ideas, above everything else, are what makes the digital domain tick. Taken further, a group of children from a socially and economically excluded area on the outskirts of the city were enthusiastically encouraged to create an animation that exploited both traditional and digital techniques to great effect.

The following examples elaborate on these projects.

Study 1: The City Car for the Future

This project involved collaboration among the DDS, the Ford Motor Company, SGI, and Wylies, the largest Ford dealership in Scotland. It took the form of a design competition for primary schools in the Glasgow area. The brief required each student to create a design for an economical and ergonomic city car for the new millennium. A short list was selected for final judging, and the winning designs were selected by a panel including Tom Scott, Ford’s leading design guru in Detroit (and GSA Alumni).

Twelve-year-old pupil Craig Vesey created the winning design. His design concept was then developed as a virtual prototype and animated in 3D using Alias Autostudio by masters student Laurie Tate at the DDS. In addition, the 3D model was adapted for stereoscopic viewing on the studio’s large-scale Fakespace Immersion table powered by an Onyx2 Infinite Reality workstation, much to the delight of Craig and his teachers. SGI finished off the event by presenting Craig with a Nintendo 64 console complete with racing games.

The event encouraged participation not only by the school children, whose previous knowledge of 3D animation came mainly from the cinema and computer games, but also from their teachers, who gained an insight into a far wider range of applications for creative computing. If there is a particular point to be learned from this exercise, it was the surprise among some of the teachers that 3D had creative applications beyond entertainment. Part of the lesson learned must be to encourage training of teachers at all levels of education and the curriculum in the diverse applications of digital media.

Facilities such as the Digital Design Studio are ideally placed to develop collaborative projects such as this that can bring together academic and industrial partners for the benefit of both.

The City Car for the Future project led to a number of questions regarding how 3D design and animation could be taught to younger school children. The language of 3D animation is bound up with impenetrable code words such as inverse kinematics or non-uniform rational B-splines. This language is barely accessible to those of us with 10 years experience with the medium. We still have to explain to adult students why a NURB is not the singular of NURBS. So what chance do we have with the younger student?

Study2: The Tail That Wagged The Dog

When the City of Glasgow commissioned The Lighthouse gallery in 1999, the intention was to establish a permanent public display of the best examples of design in Scotland. The DDS took up the challenge to create an interactive teaching tool that attempted to demystify some of the jargon and techniques behind 3D computer animation. The prototype is on permanent display at The Lighthouse.

“The Tail That Wagged The Dog” is aimed at an audience with the literacy skills of an average 10 year old. The piece consists of a series of animated scenes in which Ralf (the dog) experiences a recurring dream that takes him on an adventure. The sleeping Ralf falls from his bed only to land on top of a speeding train. He then falls between skyscrapers, and is swallowed by a nightmare vision of his own pet goldfish, which then takes his place in bed as Ralf wakes in the goldfish bowl. The scene is animated in Maya in the style of a cartoon. Ralf is comical and something of an anti-hero.

After watching Ralf's short adventure, the viewer is introduced to three scenes from the animation. In the first scene, Ralf is clinging to the top of the train as it speeds along the track. The viewer is given a number of icons to choose from, and with a little experimentation, quickly discovers that if you click on an icon (fish, bone, or tartan ribbon) it sticks to the cursor and can be moved and dropped onto the animated scene. On the second click, the texture in the scene changes to match the icon. For example, if the bone is dropped onto Ralf himself, he changes pattern from brown fur to a pattern of bone images. Likewise, the same method could produce a tartan sky or a fish-skeleton roof. With a little further exploration, the viewer can find access to a "behind the scenes" route to simple explanations of what a shader is or how a texture is used to change the surface appearance of a 3D model.

In a similar fashion, the second scene allows the viewer to influence how Ralf reacts to a high fall from a washing line strung between two skyscrapers. Depending on which open window the viewer selects, Ralf will either bounce, crash through the ground or land with a splash. Exploration "behind the scenes" produces an explanation of dynamics and particle systems.

The third scene finds Ralf underwater (prior to being consumed by his goldfish!). Selecting from three icons provides information on how he is constructed. Illustrations and information about meshes, inverse kinematics, and other modeling and animation issues are available.

The title, "The Tail that Wagged the Dog," comes from the first simple explanation the author received of how inverse kinematics works, courtesy of Chris Ford at Alias|Wavefront, circa 1992.

The project is very much a prototype designed to keep the viewer's attention for about 15 minutes at a time, principally to avoid a bottleneck in the gallery. The intention is to expand the project to cover the animation process in greater depth, but the premise of using fun and exciting visuals to illustrate complex issues in an accessible manner will stay central to any future development. Obviously, it makes sense to use the medium of 3D computer animation to illustrate and demystify the concepts inherent in the process itself.

Feedback from the exhibition organisers and the public so far suggests that the exhibit is proving to be very popular, and not only with its target audience.

Study 3: Build Up

The final example involved collaboration and co-operation with an art gallery, a firm of architects, a children's charitable organisation, and the Digital Design Studio, no mean feat in itself. As part of the Collins Gallery touring exhibition to celebrate the anniversary of pioneering animator Norman McLaren, a new animation was commissioned to celebrate the concept of "Homes for the Future."

The Easterhouse After School Project, in an area of the city neglected by social and economic development for many years took up the challenge. A group of local children, members of the summer school, created zoetropes, direct animation, and rostrum sequences based around the theme. Supported by Arts is Magic and the Digital Design Studio, the many disparate elements were combined with specially commissioned music to create a highly successful end result. The finished animation has been screened in festivals around Scotland and is likely to be shown on British television during the year 2000.

During the production, the children (between seven and 14 years of age) visited the DDS and were introduced to some of the creative tools available on a daily basis to the school's masters students and industrial research partners. The children's response was one of open acceptance of the technology. There was no barrier, no technofear. They just wanted more information and more access. They also wanted it now! Their attitude seemed to be: Why do I have to finish school before I can learn how to do this?

Sitting as a group in the studio, surrounded by advanced (expensive!) technology, they were asked what was the most important tool that they could use in their animation project. The answer they came up with was "our imagination." They have a point, but the responsibility of the educators is to equip them with a route out of their "excluded" environment. This should include access to the best tools and teaching to express their ideas and feed their imagination.

An online gallery, Artlife, is in development. The images from the Build Up workshops and artworks produced are to be placed alongside movie clips from the animation. The site is intended to act as a resource for similar groups and their organisers as well as a point of contact for young artists around the globe.

THE FUTURE

Future plans include development of 2D and 3D workshops for the Easterhouse group and others using software freely available over the Internet.

The Teddy2 software, previewed to great acclaim at SIGGRAPH 99 in Los Angeles by Takeo Igarashi of the University of Tokyo, will be the basis of the first classes (courtesy of Takeo). The Teddy software can be used by schools, community groups, and students at no cost and is the most user friendly, gestural modeller available today.

3D animation will be introduced using Alice 3D based on the Python programming environment, again freely available via the Internet from Carnegie Mellon University. Perhaps the major 3D animation software developers will follow suit and offer free (including maintenance-free!) cut-down licenses to primary schools and community groups. You only need to look at Netscape's early strategy to find the business model.

We could have a whole, inclusive, generation of empowered visual thinkers pushing the technology forward instead of an elite minority dictating how the tools should work and who should use them. Put the tools in the hands of young thinkers and then listen when they complain that the tools don't work. The fault is not with the user. It is with the designer. The toolset will only become more accessible and inclusive when the user and designer don't have to learn a new language every time the software is "improved."

CONCLUSION

From our experience, we may conclude that with the right tools it is never too early to deliver creative digital techniques to a primary-school audience. The response from a seven-year-old pupil is almost identical to that of a 37-year-old masters student when those first images appear. It is one of astonishment and surprise, quickly followed by exploration and discovery.

Our intention is to continually develop projects that will attempt to inform and encourage creative practise with digital technology in all areas of the community.


The projects described here illustrate the balance that we are attempting to achieve at the Digital Design Studio. Where it is necessary and often economically desirable to collaborate with large multinational companies on leading-edge research projects, it is also important to remember who the technology users and consumers of the future are going to be. By the time our advanced research has filtered into the mainstream of society, the requirement for a highly skilled user base will be greatly increased. If we start to deal with this now, by adopting a policy of inclusion, we will not only push the technology farther than our current horizons, we may even get digital tools that respond to our imagination rather than the other way around.

References

1. Teddy2. www.mtl.t.u-tokyo.ac.jp/~takeo/teddy/teddy.htm or distributed with Alice 3D: www.alice3d.org
2. Glasgow School of Art: www.gsa.ac.uk
3. Furness, M. (1998) *Art in motion, animation aesthetics*. J. Libbey & Co.
4. Blair, P. (1994) *Cartoon animation*. Walter Foster Publishing.



"The Tail That Wagged The Dog"



"Build Up"

Developing a Curriculum on the Arts and Mechanics of New Media for an International Post Graduate Audience

L. Miguel Encarnaç o
Fraunhofer Center for Research
in Computer Graphics, Inc.
mencarna@crcg.edu

The International Certificate Program for New Media (ICPNM) has been developed by Fraunhofer CRCG in cooperation with the Rhode Island School of Design/Continuing Education (RISD/CE) for international, early-to-mid-career professionals in government, business, industry, or education. Students from these fields benefit from a rigorous and comprehensive course of study in a concentrated six-month period with an additional three-month period of exclusive involvement in ongoing projects. Other program partners that provide additional know-how and lecturers are Technische Universität Darmstadt, Germany; the Media Communication Management Institute at the Universität St. Gallen, Switzerland; the Fraunhofer Institute for Computer Graphics, Darmstadt, Germany; Brown University, Providence, Rhode Island, USA; and the Centre for Advanced Media Technology, Singapore.

Accompanied by Web-based course material for distance course preparation and review, the program offers a sequence of learning experiences that encompasses a specific body of knowledge and theory regarding design, use, and applications of new media. This program, however, goes beyond a theoretical exploration of new media issues by grounding the theory in practical applications. Participants spend over half of their time in an ongoing research program under the auspices of CRCG and are assigned sub-tasks based on their interests and the needs of the project. At the same time, they explore design processes with RISD faculty and are immersed in four core topical content areas: 3D modeling, animation, user interface and Web page design and development, and cross-media publishing.

In this paper we present the motivation that led to the development of the curriculum, the resulting program structure, the content and organization of the program, and the outcomes and lessons learned from the first two cycles of the program, which began in October of 1998.

INTRODUCTION AND MOTIVATION

Digital libraries, e-commerce, and online communities are topics that not only draw much interest from the business world today, but also represent markets with high potential and expectations. New media, which combines many of the activities in these areas, has recently become a common buzzword within the business community.

In the areas of information technology (IT) and especially new media, there is an increasing shortage of qualified personnel. In Europe alone, there is an estimated deficit of 1.3 million qualified workers in the IT sector.¹ In an effort to turn the tide, universities and educational institutions worldwide have begun offering more programs and courses to provide the professionals demanded by the marketplace. While courses in Web design, multimedia design, CD-ROM design, and lately new media design are generally offered at the university level, an increasing number of continuing education courses in these areas are also becoming available to the growing number of young professionals in business who are trying to stay abreast of new technological developments.²

Topics being taught include computer graphics fundamentals; Web design and programming with HTML, XML, and Flash; CD-ROM design; image and video processing based on JPEG and MPEG; and computer animation. However, most of these courses are application- and tool-specific, and do not fully address important design issues.

Despite the proliferation of these courses, professionals in the field of new media have yet to agree upon exactly what new media means and what the key aspects of it are, which leads to the question: What comprises a good curriculum for students in this area?

It is difficult to give a good definition of the term “new media.” A working definition given by the Canadian Radio-Television and Telecommunications Commission (CRTC)³ is as follows:

New media can be described as encompassing, singly or in combination, and whether interactive or not, services and products that make use of video, audio, graphics and alpha-numeric text; and involving, along with other, more traditional means of distribution, digital delivery over networks interconnected on a local or global scale.

New Media projects, however, do not deal only with commercial production and distribution. Of similar importance are profound knowledge and skills in the areas of organization and management, cultural and social structures and processes in the application domains, and interdisciplinary understanding of new media requirements and impacts.⁴ New media have a great integrating effect on working processes and structures that traditionally were separated and distinct, thus allowing for an in-depth specialization of personnel and departments. In today's industry, the demand for strictly specialized education and training is regressive. The new media industry increasingly requests well-educated professionals with hybrid educational backgrounds and a broad overview on many aspects of new media. This is especially true in the lower to middle management. Additionally, working in multidisciplinary teams is an absolute requirement in the modern workplace, but it is hardly considered by traditional curricula.⁵

Universities and other traditional educational institutions cannot adapt curricula fast enough to meet this challenge. Traditional curricula are mostly uni-disciplinary with, at most, a few interdisciplinary courses. Historically evolved competence structures within the traditional educational framework hamper the development of multidisciplinary courses and classes, which, in turn, obstruct development of team-oriented learning environments. In addition, the teaching personnel that trained in the “old media ages” are not keeping up with the rapid evolution of new media. Consequently, new media techniques and technologies are rarely employed in today's new media education.

Several attempts to overcome this dilemma have produced either courses that give an overview of new media to an interdisciplinary audience, with titles such as *New Media for Everybody* or *Everybody's New Media*, or in continuing education programs that are strongly focused on select aspects of new media (for example, visual arts, technology, or business).

These attempts do not fulfill industry's demands for qualified new media professionals as described above. A new educational paradigm is required to improve current practices and to prepare students to work in an international, multidisciplinary, and team-oriented context.

THE NEED FOR A CROSS-DISCIPLINARY EFFORT

In order to meet the requirements of today's new media industry and anticipate those of the future, we identified the following needs for a corresponding post-graduate curriculum:

- Communicate a holistic picture of new media. New media is more than a single discipline; it is a working style, even a lifestyle, that incorporates many different traditional disciplines, is influenced by them, and vice-versa. Key disciplines for new media include the arts (aesthetics, human-centered, and creative aspects), technology (new media hardware, software, networking requirements, and developments), and business (economic impacts of the new media revolution).
- Teaching and learning in a multicultural setting. Due to its unique role in today's industry and society, new media is strongly influenced by the cultural context in which it is employed. Cultural context can be geographic, social, or professional. We incorporated multiculturalism into the ICPNM program by selecting an international student audience with different professional backgrounds and by involving an international and multidisciplinary team of teachers from universities, industry, and training institutions. Furthermore, we taught courses in the facilities of different participating educational institutions that have different teaching cultures.
- Teaching in multimedia settings. New media education is not convincing if the topics are not incorporated and used in the education delivery itself. We therefore employ state-of-the-art multimedia classroom, computer, and tele-conferencing equipment, as well as Web-based course material for preparation and review, to make the course a new media experience.
- Training in multidisciplinary teams. We consider teamwork a key component of new media working environments. Therefore, special emphasis has been placed on developing the corresponding skills of the course participants. Small, continuously changing, multidisciplinary teams are created to fulfill common tasks, while the group as a whole is charged to complete larger projects. By design, these projects cannot be completed by a single participant. Special team-organization seminars designed for this program prepare the participants to work on teams and as teams.

To meet all the requirements described above, the participating educational institutions had to be carefully chosen in order to provide maximum flexibility of curriculum development while maintaining the highest educational quality. In addition, the different key aspects of new media as described above had to be represented, ideally in an educational context. Choosing Fraunhofer CRG as the incubator for cooperation between the Technische Universität Darmstadt, the Rhode Island School of Design, and the Media and Communication Management Institute of Universität St. Gallen, created a collaboration of leading international educational institutions with industrial affiliations in new media technology, the arts, and business. This collaboration is supported by the INI-GraphicsNet, an international network of excellence comprised of institutions for advanced education, training, and research and development in computer graphics technology, systems, and applications, and represented by Fraunhofer CRG.

The initial challenge in developing an integrated curriculum was to find a common understanding and terminology for similar new media components and different approaches; language barriers did not arise because of different nationalities, but rather because of the participants' various professions and working practices. Commonalties had to be carefully scrutinized and differences thoroughly studied before the educational new media topics described below could be identified.

IDENTIFIED EDUCATIONAL TOPICS

Based in part on the specialties of the participating educational institutions, the overall new media topics were identified: computer graphics, visual arts, multimedia, cross-media publishing, telecommunication, e-commerce, media management, and multimedia marketing. Each of the educational institutions contributes to one or more of these topics. The main media upon which the program focuses are image, audio, video, Web, and CD-ROM. Depending on the program emphasis, these overall topics lead to different courses to be taught. The next step, therefore, was to decide on the emphasis of the pilot program.

Choosing a Pilot

We strongly believe that the main benefit of new media is the support it provides for communication and interaction between people, collocated or geographically dispersed. The main components of new media communication in this context are, therefore, humans, computers, and networks. Consequently, we decided to choose a pilot program that focuses on supporting interactivity between humans and the computer medium, as well as between humans and humans through computers, networks, and other media. Using "principles of interactivity" as the program's focus allows us to cover such diverse areas as screen and user-interface design, human-computer interaction, telecommunication, virtual reality and e-commerce.

Developing a Curriculum on The Arts and Mechanics of New Media for an International Post Graduate Audience

Developing the Curriculum

The ICPNM program is split into five modules, each of which last between three and six weeks. Modules consist of a combination of topical courses, workshops, lectures, and seminars. Each module begins with an introductory course on computer applications and is complemented by a larger project that runs throughout the module. Courses last about one week and include approximately 20 contact hours and additional lab times. Technologically oriented courses are taught by Fraunhofer CRCG, and courses with visual arts emphasis are taught by RISD/CE. They are thematically synchronized by the schedule and integrated content-wise by the larger module-spanning projects. In each of the courses, smaller exercises allow for preparative work on the main project.

Workshops are whole courses that put an emphasis on lab practices. Business-related lecture series, which accompany the instructional part of the program, build upon themselves, thus completing the other program modules as a separate track. The program also provides separate one-time guest lectures or colloquia in the form of seminars, on topics such as teamwork or the state-of-the-art of interactive 3D graphics.

A detailed version of the course description can be found in the supplementary materials to this publication on the corresponding SIGGRAPH 2000 CD-ROM.

LAYING OUT A MAP

In order to structure the curriculum, a mapping had to be found between knowledge and skills to be taught and content areas to be covered. First, content and structure were defined in detail for each course. For each subtopic of a course, knowledge and skill outcomes were listed and rated. Here, knowledge and skill outcomes were structured in the areas of:

- Design principles: color, composition, contrast, and form.
- Interactive principles: time, audio, video, and navigation.
- Technical principles: mathematics, algorithms, and security.
- Conceptual principles: creativity, narrative, and media.
- Behavioral principles: communication, management, and teamwork.
- Computer applications: Photoshop, Director, Maya, etc.

This complete map (see Figure 2) was used to tune the program and course structures by identifying overlapping topics and missing lectures, and by adjusting the course order in the program. This structure was also used to identify lecturers who might fit into the program.




Figure 1: ICPNM program structure overview.

Topic	Design Principles				Interactive Principles				Technical Principles			
	Color	Form	Contrast	Composition	Time	Audio	Video	Navigation	Mathematics	Algorithms	Security	Other
Cognition	*	*	*	*								*
Content	*	*	*	*								*
Color	*	*	*	*								*
Typeface	*	*	*	*								*
Form	*	*	*	*								*
Balance	*	*	*	*								*
Processes	*	*	*	*								*
Structure and Hierarchy	*	*	*	*								*
Fundamental Language Terms												*
Shapes in 2D/3D												*
Complex Interfaces												*
Action/Reactions												*
Animation in 2D/3D												*
Interaction in 2D/3D												*
Building the 2D/3D Layout												*
Building the 2D/3D Layout												*
Building the 2D/3D Layout												*
Building the 2D/3D Layout												*
Building the 2D/3D Layout												*
Applying Today's Software												*
Processes												*

Figure 2: Section from a table mapping targeted outcomes to program modules and assignments.

MAPPING CONTENT AND SKILLS TO ORGANIZATIONAL STRUCTURES

In order to organize and schedule classes and other program modules, the intersection points in the map had to be converted into course modules by combining similar, complementary, and coherent modules from the participating institutions that would cover the corresponding, knowledge, skill, and content requirements. The results led to the curriculum described above. The order of the resulting modules was not only a scheduling issue but also was based on the program structure, described in the following, that aimed for an effective learning experience.

STRUCTURING THE PROGRAM

The program is built around four essential phases. The first three phases occur for participants in cycles throughout the program. All four phases, however, are also distinctly separate in the structure of the program:

1. Acclimation

During this opening phase, program participants receive an orientation in three cultures: the larger American culture, the specific research culture of Fraunhofer CRCG, and the art and design culture at RISD. Participants explore these cultural contexts through activities, orientation sessions, site visits, discussions, and role-playing. During this time, they also examine student/teacher, supervisor/supervisee, coach/mentor, and team member roles. Furthermore, students explore cross-cultural perspectives and familiarize themselves with the specialized languages and processes of art, design, and technology.

2. Immersion

During the second phase of the program, participants immerse themselves in their preferred content areas in order to explore the processes of design and the applications of new media. Each student focuses on four content areas and explores their applications through a practicum. Students are expected to divide their time between RISD studios and CRCG labs, where they engage in structured learning activities in topical content areas, such as 3D modeling, animation, user interface development, Web-page design, and cross-media publishing. Fraunhofer CRCG provides open lab hours for practice, project assignments, and practicum work. Students work with RISD faculty on both the design process and the designed product across different media, while CRCG staff and supervisors provide project development and delivery through the practicum experience.

3. Reflection

The third phase of the program is a formally structured opportunity for participants to step back, review, and articulate learning activities by applying their new skills. Students are asked to document and present a digital portfolio, which communicates something about their experience and what they have learned. This portfolio could take the form of a Web page, multimedia presentation on a CD, or a more traditional form, such as a published book. During this reflection time, students contemplate and critically assess the processes and product outcomes of their time spent in the program thus far. Their work is reviewed by both RISD faculty members and CRCG professionals, who provide feedback on the presentation's effectiveness as a communication vehicle and its success as a summative and evaluative tool.

PRACTICUM

In this final phase of the program, students are assigned to ongoing projects at CRCG based on their skills, interests, and the topical content areas upon which they focused during the first three phases of the program. During this three-month period, program participants are assigned subtasks based on their interests and each project's needs. Subtasks can take various forms, such as: integrating a new 3D interaction device for immersion in virtual environments, redesigning an application system's user interface based on user-centered aspects of human-computer interaction, designing and developing a CD-ROM for an industry partner's advertising, or developing and implementing interactive Java applets to enrich Web-based courseware for computer graphics or graphic design education.

Sample project activities, which provide participants with hands-on experience, include tele-collaboration tools for e-commerce, modeling historic buildings for cultural heritage preservation, cross-media information production and delivery, and digital security. By teaming up with CRCG professionals on these projects, students may actively participate in both the communication process and the team learning experience in a supportive, professional environment.

Although the work is supervised, guided, and reviewed by CRCG professionals, the students have the freedom to develop, present, and implement their own ideas, thus emphasizing development of skills in the areas of technology evaluation, integration, and design as a process. To demonstrate their understanding of these processing skills, students are asked to document and present their project tasks, background, solution approach, and results in a report that matches academic requirements.

Developing a Curriculum on The Arts and Mechanics of New Media for an International Post Graduate Audience



Figure 3: New ICPNM Web site designed in a student project www.icpnm.org/Course98

RUNNING THE PILOT

Designed as a pilot program, the 1998 International Certificate Program for New Media (cf. Description in Part One) focused on selected international post-graduate students who benefited from a rigorous and comprehensive course of study in a concentrated six-month period, with an additional three-month period of exclusive involvement in ongoing projects.

The pilot program started on October 5, 1998 and provided education and practical experience to a class of nine selected students from India, Malaysia, Korea, and Germany with a masters or comparable degree in the areas of information technology, architecture, graphic design, or engineering. All students successfully completed this program. After the first five modules, the students had to define and run a project. Examples of such projects are redesign of the ICPNM Web site (see Figure 3), an interactive 3D walkthrough of historic settings in Providence (see Figure 4), and a Flash Animation illustrating the purpose and execution of digital watermarking agents (see Figure 5).

The current course began on October 4, 1999 and consists of 10 participants from Singapore, Germany, the US, and Korea with students who have a BS or comparable or higher degree in the areas of information technology, architecture, graphic design, economics, and philosophy.



Figure 4: A snapshot from the animation 'Virtual South Main Street, Providence, RI'.



Figure 5: Data I/O depiction in the Digital Watermarking Agent Animation.

EVALUATION AND MODIFICATIONS

A direct demonstration of the success of the program is the fact that during the course of the program all students found jobs in new media, covering the whole area of possible working locations from Internet companies to ad agencies to research institutes.

To gather detailed feedback on the program, we performed an evaluation together with the students. Students were asked to rate the individual courses and to describe their experiences in the program. The primary criticism addressed the initial lack of coordination between assignments within individual courses. Students preferred assignments on individual topics, such as image processing, where the results are consequently reused in subsequent courses and modules, thus giving them a more complete view of the production and management processes. Based on the experiences from the pilot and faculty evaluation meetings during the first run of the program, a modified and more structured assignment schedule was developed for the second run.

EXPERIENCE AND LESSONS LEARNED

The diversity of the program participants (the different professional backgrounds, credentials, and expectations of students and instructors) made ICPNM a learning experience for everybody, including the organizers. Since the outcomes could hardly be predicted, the course was managed with a great amount of flexibility, allowing for integration of participants' input and propositions. The following lessons learned from the instructors and organizers are exemplary for the dynamics of the program:

- Don't underestimate the students' learning abilities! It was amazing to experience how rapidly students picked up knowledge and skills from areas that were completely outside their own experiences. The main reason for the success was that students were suddenly confronted with information that was complementary to their own experiences, that they "always wanted to know." Furthermore, working in teams with other students gave them additional motivation to demonstrate flexibility.

- Don't underestimate the students' motivation! The number of contact hours required in the program and the complexity of the final projects presented a stamina challenge to the students. Nevertheless, their craving for knowledge never decreased, which challenged the lecturers to continuously add, modify, or deepen presented materials and to introduce additional exercises.
- Don't underestimate the cost of keeping up with the state-of-the-art! In a program that claims to teach state-of-the-art new media skills and applications, the newest of the new is absolutely required. Previous software versions, last year's computer hardware, or traditional classroom equipment were categorically criticized by the students, who tried to "max out" any available application or system.
- Don't underestimate the effort to integrate a multidisciplinary curriculum! In addition to finding a common terminology, as mentioned before, many issues arose during development of the curriculum that one must anticipate to avoid drowning in administrative overload: different educational environments have different processes with different timelines; different business and budgeting models must be coordinated; competencies need to be clarified; guidelines for teaching must be distributed; and a general awareness of the main goals and intentions of the program must be established.
- Don't underestimate the effort to organize an international group of teachers! International experts are not only rarely available for teaching, but their schedules also constantly change. Scheduling these experts early, constantly updating their visit dates, and revising the program schedule around these changes are required to make the program run smoothly. In addition, one cannot expect the different instructors to communicate with each other in order to synchronize similar or complementary courses; instead, this must be done by the coordinators in cooperation with the instructors.
- Don't overestimate the appeal of distance learning! The ICPNM program offered different participation models ranging from remote course participation for the practicum to partial Web-based participation during the first six months of instruction. Although this option reduced the course cost significantly, no participants shortened the nine-month stay in the US in favor of the lower cost or to reduce their time away from the workplace.

CONCLUSIONS AND FUTURE ACTIVITIES

The first two runs of the ICPNM program were great successes. The program has been extremely well received and our own promotional efforts are increasingly supported by personal recommendations from former participants. International acceptance has stimulated strong interest in the program by other educational institutions that wish to join the program. However, this must be carefully evaluated, since the current program is very successful in its current state, and the organizational effort is already high. On the other hand, we are discussing possible new business models (for example, how to incorporate ICPNM into other university and MBA curricula in the form of intensive study abroad). For this reason, we will continue our efforts to support remote participation in ICPNM, which will strongly depend on development of Web-based approaches for skill-based education. In addition, we are seeking certification of the program by appropriate regulatory authorities.

ACKNOWLEDGEMENTS

We would like to thank RISD president Roger Mandle and Fraunhofer CRCG president José L. Encarnação, as well as Edmund J. Ferszt, Joe Quakenbusch, Joanna Roux, and Matthias Wloka for their enthusiasm and support in developing the ICPNM program, and all the many lecturers from all over the world for their contributions to the program.

References


1. Ewing, J. & Dawley, H. (1999, Dec. 27). The missing worker. *Business Week*, 18-19.
2. Weber, J. (1999, Oct. 4) School is never out. *Business Week*, 98-101.
3. Canadian Radio-Television and Telecommunications Commission (CRTC). www.newmedia-forum.net.
4. Gonzalez, R., Cranitch, G., & Jo, J. (2000). Academic directions of multimedia education. *Communications of the ACM*, 43 (1), 89.
5. Keskeys, D. (1998). Media technologies and an interdisciplinary approach to program design. *SIGGRAPH 98 Conference Abstracts and Applications*, 52-59.

How do you explain difficult concepts to students? By drawing pictures? By drawing on analogies and metaphors? By acting out scenarios? These are all methods that enhance the student's classroom experience, but what about distance learning? Animation can combine these methods and help teach concepts that are difficult to describe in words and still images alone. And it can work well for students in a classroom or online.

Our team creates short animation pieces that visualize technical concepts. Education psychology research shows that students learn and retrieve information best when they can construct more than one pathway in their memory. Stimulating multiple senses when teaching new concepts helps the student create these pathways. Recall is enhanced when information is presented in both visual and verbal form. Animation combines visual symbols with an auditory narration. Students may also perceive animation as a motivator to learn challenging new concepts.

We follow a written process to work through an entire animation project. We call it the "10-40-40-10" process because its steps are 10 percent planning, 40 percent design, 40 percent implementation, and 10 percent test. First, we determine if a project is a good candidate for animation. Next, we plan the project by analyzing the audience, outlining the content, and identifying technology constraints. After the plan is approved, we write scripts and design the animation by drawing storyboards. With designs in hand, we implement the animation. We review the animation as we work on it, and make adjustments as necessary.

Finally, we deliver the animation to the audience, either on the Web, or as part of a CD-ROM multimedia presentation. Animation takes time to develop, but we have found it to be well worth the effort. People who have viewed our animations tell us they enjoy them and gain a better understanding of the concepts we are presenting. Multimedia tools are now more available and easier to use, making it possible for anyone to enhance students' learning with simple animation.



IBM uses animation to illustrate concepts that are difficult to describe in words and still images. We follow the "10-40-40-10" process to plan, design, implement, and test animated presentations.

INTRODUCTION

Immersive Virtual Reality (IVR)¹ has exciting potential for educational applications, because it can enable students to interact with ideas in new ways. In an IVR environment, students can engage in simulations of real-world environments that are not accessible due to financial or time constraints (high-end chemistry laboratories, for instance), or which are simply not possible to experience, such as the inside of a volcano or the inside of an atom. IVRs also enable students to interact with visualizations of abstract ideas (for example, mathematical equations and elements of color theory). The hands-on, investigative learning most natural to IVR is an excellent way to train new scientists and engineers. In addition, because the environment is computer-generated, it is an ideal future platform for individual and collaborative distance learning efforts.

IVR lends itself naturally to a constructivist pedagogical approach, one that builds on Piaget's view that knowledge is created by learners through interaction with people and the world around them⁷ and that it is characterized by first-person, non-symbolically mediated interactions.²⁰ But the success of such virtual "direct experience" depends heavily on the design of interface and interaction techniques. Although IVR may suggest ease of use and obvious, transparent interface designs, the opposite is, unfortunately, true: IVR presents surprisingly difficult interface challenges. Even techniques as seemingly straightforward as object selection and navigation are the subject of dozens of research papers.^{13,14} The design of IVR interaction that feels natural and promotes learning is an almost entirely unsolved problem. We have been tackling some of the interface and interaction challenges of using IVR for education in a Museum of Color in which students can perform a number of science museum-like experiments.

Although IVR environments are still too expensive and fragile for mainstream educational use, the use of VR in training is well-established,¹⁰ and its use in entertainment is growing rapidly.^{21,1} Moore's Law is seen in the use of commodity 3D graphics in games and edutainment (the Sony Playstation 2 has more graphics power than most of today's high-end workstations), and suggests that today's esoteric immersive environments will become commodity classroom fixtures in the future. We feel, therefore, that it is critical to conduct educational IVR interface and interaction research today so that educators can take full advantage of IVRs in the near- and long-term future.

PREVIOUS WORK

The IVR Museum of Color extends ongoing work done in the Exploratories Project at Brown University (www.cs.brown.edu/exploratory). Exploratories are combinations of science museum-like exploratoriums and laboratories, realized as 2D and 3D explorable worlds. The Exploratories Project has a dual mission: creating innovative educational content for the Web and documenting the experiences of creating such software in a design strategy handbook.

A handbook for 2D and 3D desktop applications is already in progress as part of the Exploratory Project.^{8,19} The handbook draws on our own and others' experiences and includes guides, templates, patterns (as in Gamma et al's object-oriented "Design Patterns"), and examples. Topics include assessment of resources, descriptions of different pedagogical categories, patterns for implementation, indexed examples, and, of course, extensive information on interface and interaction design.

The Exploratories project has previously focused on 2D and 3D graphics on the desktop, creating, for example, suites of 2D applets for teaching color theory^{2,17} and basic signal processing used in image processing. The project has also worked with 3D desktop graphics, creating modules, for example, for teaching synthetic camera transformations. Recently, we have begun to design for semi- and fully immersive environments, from the Barco Baron table (a rear-projection display) to a four-walled CAVE. Building on the two suites of 2D color theory exploratories we have created, we have begun creating a sequence of CAVE-based exploratories, a Museum of Color.

IVR education projects at a number of universities have already shown promising results: for example, the University of Illinois Electronic Visualization Laboratory's Round Earth project,¹² George Mason University and the University of Houston Virtual Environments Technology Laboratory's NewtonWorld and MaxwellWorld,¹⁶ the University of Michigan Virtual Reality in Chemical Engineering Laboratory's Vicher (Virtual Chemical Reaction) modules 1 and 2,³ the University of Washington Human Interface Technology Laboratory's Zengo Sayu,¹⁵ and MIT's ALIVE environment,¹¹ just to name a few. We have been convinced, by reading about these works and seeing demonstrations, that the educational potential for IVR is real, but we have found little specific discussion of interaction issues that are particularly relevant to educational uses of IVR. By conducting projects like the one discussed in this paper and presenting interaction methods that we find successful, we hope to make construction of future educational IVRs easier for ourselves and others.

ENTER THE MUSEUM

The Exploratories Project has begun work on an IVR-based Museum of Color to teach difficult-to-understand, often highly abstract aspects of color theory through a series of hands-on experiments that take advantage of immersion in a 3D virtual space and interaction with explanations of abstract concepts. This project has illuminated a number of issues in IVR-based interaction relevant to instructional design; we discuss three of them here.

1. An IVR is created with a combination of immersion (3D graphics rendering that surrounds the user, with the point of view dynamically updated by tracking the user's head position and orientation), and the use of interaction devices such as data gloves. A 3D world seen on a desktop machine monitor does not qualify as an IVR, even when stereo is used, because the sensation of immersion is absent.

1. Using Architectural Spaces to Structure a Sequence of IVR Learning Modules

One of the first challenges we encountered was how to navigate through a complex series of learning experiences in an immersive environment. In our desktop exploratories, we use customized hypertext structures that we have developed for presenting educational material on the Web.³⁸ For example, we have structure templates for games, lab-type experiments, and more free-form playground environments. In the Museum of Color, we chose an architectural metaphor that guides the user through various learning sequences and simplifies nonlinear navigation. The 3D structure also serves as a mental model of the lesson structure. While we have not found this type of architectural interface to be terribly successful for the desktop, we have found it to be quite natural in the CAVE. In the Museum of Color, for instance, each floor has the same layout and flow sequence (see Figure 1) and higher-numbered floors house experiments of increasing conceptual difficulty. Simple interactive experiences with different parameters of a color (such as value, hue, and saturation) are on the first floor, experiments dealing with groups of color are on the second, and interactive comparisons of color spaces are on the third. When a user has completed a floor, the architectural metaphor, complete with an elevator ride between floors, enhances both the sense of having traversed a memorable territory and the sense of closure over a set of concepts.




Figure 1: Flying over the floor plan. On each floor, a central column (which holds the elevator) is surrounded by four rooms.

2. Trade-offs Between Metaphor Integrity and Convenience

We have found that the effectiveness of using an architectural space depends first on its appropriateness to the task, and secondly on the correct balance of spatial integrity and navigational convenience. Faithfulness to the metaphor is essential for establishing it as a mental model. However, undue faithfulness can be irritating. For example, in the Museum of Color, users can take an elevator between floors to emphasize the structure, but we also let users navigate in a nonlinear fashion by selecting rooms on a 2D map. The step outside the 3D realm for map navigation is similar to use of maps in real life. We avoided flying, passing through walls, and jumping or transporting from place to place within the museum. For educators, establishment of the structure can be further exploited by the equivalent of hypertext guard fields.⁵ Users can be forced to follow a specific sequence and can be barred from certain rooms until specific conditions have been met, such as conducting a number of prerequisite experiments. Our framework allows for such guiding.

Sometimes the benefits of one approach, such as use of an architectural metaphor, clash with another potential interface or interaction scheme. For example, scale changes are a particularly powerful feature of IVR learning environments, but it is difficult to take full advantage of the potential for relative scaling while still maintaining the visitor-in-a-museum metaphor. In our color-space rooms, users see a number of different 3D color spaces, rendered as wireframes, placed on pedestals like sculptures (see Figure 2). The color spaces can be investigated in a number of ways. A series of visualization techniques helps to show off the internal structure of the spaces (see Figures 2, 3). Two-handed interaction (using data gloves) lets users pick up color spaces and rotate them. Before they were integrated into the museum, it seemed natural to scale the spaces and walk around inside them. Inside the museum, such radical scaling causes the color spaces to become larger than the room and intersect with the walls and ceiling. A partial solution is to restrict the ultimate scale; another approach is to let users “jump” into the color space (See Figure 2).




Figure 2: An HSLcylinder, rendered as a wireframe, with falling disks colored according to their changing locations. The user approaches (left) and then enters the space (right).



Figure 3: Ribbons help to convey the location of colors in an HSL cylinder without obstructing the user's views of the space.

3. Concept Visualization

An interaction area important for making IVR useful for the widest possible range of educational topics is the use of 3D representation (and thus manipulability) of parameters not usually visualized in 3D, an approach that we call concept visualization, similar to Winn's "reification."²⁰ The first room of the Museum of Color teaches the concepts of value and saturation with this method. Traditionally, users manipulate a value or saturation parameter with a slider or dial, or by moving a brush-type tool over an area of their image. In the CAVE, we present the changes to either parameter in an image as a height field (with the image mapped onto it) that users can drag and push with gloves to get a sense of what happens when you shift either parameter of a color (see Figure 4). Interacting with a 3D representation of value or saturation gives the user a new mental model of its properties and lets manipulation draw on use of the fingers, hand, arm, and even full body. The 3D representations of these abstract parameters will, we hope, give them a more tangible reality for students and continue to enrich artists' and designers' vocabulary for talking about the use of value. For example, artists already often describe black as heavy and white as light or refer to light or dark areas as receding or advancing. Interestingly, the model of saturation and value as height appears to persist in users outside of the CAVE, and anecdotal reports show that it helps users create and analyze 2D works in the real world.



Figure 4: A researcher uses a pinching motion to warp the height field, in this case making a pale blue area more saturated by drawing it upward.

Visualization of color spaces is more direct. Usually, three parameters are involved, such as hue, saturation, and value or red, green, and blue, and the visualization of each parameter's values along an axis serves to make these somewhat arbitrary parameters more tangible. For beginning students of color theory, it is often not clear why one needs so many different color spaces and what their different purposes are. One of the most important distinctions is between spaces that are perceptually uniform (an equal-distance step along any axis results in an equal amount of perceived change in the quantity of the axis's parameter) and those that are not. We are able to let users compare perceptually and non-perceptually based spaces by using a fully manipulable cutting plane that is flat in a perceptually based space but maps to a warped shape in spaces that do not have perceptually linear changes along the axes of their various parameters. Although one can do this on the desktop, we found the immersive experience to be qualitatively different. The ability to step inside a color space, move the plane with one's arms, and easily examine the results from any point of view gave us a much better intuitive sense of shapes of the spaces and the areas in which perceptual changes become more and more nonlinear.

CONCLUSION AND FUTURE WORK

The interface challenges discussed here are only a few of the many issues confronting educators who want to tap the enormous potential of IVRs as learning environments. In addition to work being done on basic IVR UI techniques, such as those needed for navigation, selection, and multimodal interaction (for example with gesture and voice recognition),^{4,6,9} we hope that particular attention will be paid by educators to IVR interaction support for educational needs such as note-taking or record-taking, narrative structures, scenes that change adaptively to respond to a learner's growing sophistication, tracking and remediation of user mistakes or misconceptions, collaborative learning and problem solving, and performance measurement.

We plan to continue to add content to our Museum of Color, and to create further educational IVR experiences. But our long-term goal is to document both our experience and that of others to codify solutions to some of the many problems that arise when using IVR for educational purposes. Although a number of successful IVR education projects have been completed (see Previous Work, above), a source for clear guidelines and examples will, we believe, do a greater service to the field than simply creating more content. We plan to write an IVR section for our design strategy handbook that will contain discussions of specific issues, like those presented in this paper, as well as lists and examples of the relevant IVR techniques found in today's research literature. Widespread dissemination of practical results in the handbook will make our work available to a broad range of potential end users, from research scientists to high school educators.

Our thanks to Andy van Dam for his critical reading and to David Karelitz and Jean Laleuf for feedback during the creation of the Museum. We would like to thank Adobe Systems, Inc. for their generous support of this work. We would also like to thank the other sponsors of the Brown University Computer Graphics Group and the CAVE at Brown: Advanced Network & Services, the Technology Center for Advanced Scientific Computing and Visualization, the Department of Energy, IBM, Intel, Microsoft, the NSF Science and Technology Center for Computer Graphics and Scientific Visualization, and Taco, Inc.

References

1. vr-atlantis.com/lbe_guide/lbe_list3.html
2. Beall J., Doppelt A., & Hughes, J. (1996). Developing an interactive illustration: Using Java and the Web to make it worthwhile. *Proceedings of 3D and Multimedia on the Internet, WWW and Networks*. Bradford, UK: National Museum of Photography, Film & Television.
3. Bell, J. & Fogler, H. (1995). The investigation and application of virtual reality as an educational tool. In *Proceedings of the American Society for Engineering Education*.
4. Benford, S., Taylor, I., Brailsford, D., Koleva, B., Craven, M., Fraser, M., Reynard, G., & Greenhalgh, C. (2000). Three-dimensional visualisation of the World Wide Web. In *ACM Computing Surveys Symposium on Hypertext 2000*.
5. Bernstein, M. *Storyspace: Hypertext writing environment*. Watertown, Massachusetts: Eastgate Systems Press.
6. Chen, C. & Czerwinski, M. (1998). From latent semantics to spatial hypertext - An integrated approach. In *Proceedings of Hypertext '98, ACM SigLink*.
7. Fosnot, C, ed. (1996). *Constructivism: Theory, perspectives, and practice*. New York: Teachers College Press.
8. Gould, D, Simpson, R., & van Dam, A. (1999). Granularity in the design of interactive illustrations. In *Proceedings of SigCSE '99, ACM SigCSE*.
9. Lindeman, R., Sibert, J., & Hahn, J. (1999). Towards usable VR: An empirical study of user interfaces for immersive virtual environments. In *Proceedings of SIGCHI '99, ACM SIGCHI*.
10. Loftin, R. & Kenney, P. (1994). Virtual environments in training: NASA's Hubble Space Telescope mission. In *Interservice/Industry Training Systems & Education Conference*. www.vetl.uh.edu/Aerospaceov.html; www.vetl.uh.edu/Hubble/longpaper.html
11. Maes, P. (1995). Artificial life meets entertainment: Lifelike autonomous agents. in *Communication of the Association for Computing Machinery*, 38 (11).
12. Moher, T., Johnson, A., Ohlsson, S., & Gillingham, M. (1999). Bridging strategies for VR-based learning. In *The Proceedings of CHI 99*, 536-543. www.evl.uic.edu/roundearth/
13. Pierce, J., Forsberg, A., Conway, M., Hong, S., and Zeleznik, R. (1997). Image plane interaction techniques. In *3D Immersive Environments, 1997 Symposium on Interactive 3D Graphics*, 39-43.
14. Pierce, J., Stearns, B., & Pausch, R. (1999). Two handed manipulation of Voodoo dolls in virtual environments. *1999 Symposium on Interactive 3D Graphics*, 141-145.
15. Rose, H. & Billingham, M. (1995). Zengo Sayu: An immersive educational environment for learning Japanese. *University of Washington, Human Interface Technology Laboratory, Report No. r-95-4*.
16. Salzman, J., Dede, C. & Loftin, B. (1995). Learner-centered design of sonorily immersive microworlds using a virtual reality interface. In *Proceedings of the Seventh International Conference on Artificial Intelligence and Education*, 554-564. vetl.uh.edu/ScienceSpace/learnvir.html; www.vetl.uh.edu/ScienceSpace/absvir.html
17. Simpson, R., Spalter, A., & van Dam, A. (1999). Exploratories: An educational strategy for the 21st century. In *Electronic Schoolhouse, ACM SIGGRAPH 99*.
18. Spalter, A. & Simpson, R. (2000). Reusable hypertext structures for distance and JTLearning. Accepted for *ACM Hypertext 2000*.
19. Spalter, A. & Simpson, R. (2000). Integrating interactive computer-based learning experiences into established curricula. Submitted to *ITICSE 2000*.
20. Winn, W. (1993). A conceptual basis for educational applications of virtual reality. *University of Washington, Human Interface Technology Laboratory, Report No. TR-93-9*.
21. Zyda, M. & Sheehan, J. (Eds). (1997). *Modeling and simulation: Linking entertainment and defense*. National Research Council, National Academy Press.

Is There Any There There? Design and Implementation of Successful Team-Oriented Educational Environments for Computer Animators

Organizer
Barbara Mones
University of Washington
mones@cs.washington.edu

The study and practice of computer graphics, animation, and immersive design require continued acquisition of interdisciplinary and cross-disciplinary knowledge in combination with collaborative working habits and skills. Many universities have struggled with the problem of locating courses and/or programs in appropriate schools or departments and hiring faculty with wildly diverse backgrounds within the traditional academic structure in order to best educate students with these interests.

Two models often emerge: a piecemeal approach with curriculum located in several different and often inappropriate departments, or a singular approach that selects one department to the exclusion of other appropriate ones. Both models have serious downsides. The first often results in faculty members, courses, and programs that have no information about the others, and everybody often suffers as a result. The second model can tend to shut out departments with less political clout and small budget power. The result can be an unbalanced curriculum and undereducated students who don't have the broad perspective required to succeed in the graphics and animation industry.

In particular, interaction and integration between art and computer science departments has been particularly difficult and often unsuccessful. There are programs that have been designed from the outset to serve students from many backgrounds and interests, and some have emerged over the past five years. They are actively marketed as integrated programs composed of faculty from various backgrounds and departments.

The computer animation industry has almost universally provided environments for painters, sculptors, computer scientists, engineers, and many others to cooperatively and collaboratively work in teams, because they must do so to develop quality products, whether the end result is feature films, interactive games, or immersive environments.

This panel focuses on ways to incorporate this cooperative and collaborative approach into our academic environments, including ways in which this is currently happening. Specifically, the panel looks at schools that are now providing classes that require the whole class to work as an animation team, based on the students specific expertise and interest, to produce short computer-animated films or VR environment. Students are selected because of the talent they bring to the group (which must include music, architecture, theater, computer science, fine art, etc.). The idea is to simulate the industry experience and give students the experience of working with colleagues of differing educational backgrounds and interests to meet deadlines and a collective vision as they contribute their own unique and creative input. Students may have already developed sophisticated individual approaches, but in this program, they must fit into the team and contribute in the most positive fashion.

These types of courses, when taught by experienced instructors (ideally several of them per course), offer a unique experience for students and faculty. And they provide a sense of community for the participants. Panelists discuss the process of hiring for these programs, successful methods for developing a sense of community among the students, and the specific curriculum required.

In presenting student-produced works, some that have had wide recognition and won several prizes nationally and internationally, the panelists show the process of developing a short film and cover technical and aesthetic issues, including story development and shot breakdowns. They also review the limitations and many problems, as well as special advantages, of trying to design for and build collaborative design teams in academic settings. The panel culminates with a look toward the future, with an emphasis on new technical tools and evolving approaches to animation and education.

PANELISTS

Barbara Mones was the founding director of a graduate program in computer graphics and animation called Visual Information Technologies at George Mason University. She served as a tenured professor there in the Art and Art History Department. Her work in industry has ranged from designing and developing training curricula to designing Web graphics, writing technical training materials, and animating. She has worked at NASA Goddard, Pacific Data Images, and Industrial Light + Magic. She has been a member of the SIGGRAPH Education Committee for many years and was the Panels Chair for SIGGRAPH 97. She earned a BFA at the University of Michigan, an MFA from Rhode Island School of Design and an animation certificate from Sheridan College in Ontario, Canada. Currently, she is developing curriculum and teaching computer animation courses, utilizing a collaborative approach to curriculum, in the computer science and engineering department at the University of Washington.

James Buckhouse has an undergraduate degree from Brown University in fine art. He worked at Pacific Data Images as a layout animator prior to his current position as visiting artist at Stanford University. He has created computer animations for feature film as well as fine art and experimental animation.

Gregory Niemeyer graduated from Stanford in 1997 with an MFA, and he earned a BFA in photography from Ecole d'Arts Applique, Switzerland in 1990. He founded and has directed SUDAC: CG Animation Curriculum Development since 1997. He is a practicing artist and has exhibited his work internationally. He teaches virtual object design and computer animation utilizing collaborative team approaches.

Randy Pausch is an associate professor of computer science, human-computer interaction and design at Carnegie Mellon University, where he is the co-director of the Entertainment Technology Center. He has been 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, particularly for the DisneyQuest virtual reality-based theme park. He is the author or co-author of five books and over 50 reviewed journal and conference proceedings articles. His primary interests are human-computer interaction, entertainment technology, and undergraduate education. He has taught several courses utilizing collaborative team approaches.

The Digital Design Studio of the Glasgow School of Art was developed in response to rapid advances in leading-edge digital technologies. This higher education program encompasses three distinct constituents: education, research, and training.

In general terms, it is difficult for any education system to respond rapidly to developments in advanced technology. Inevitably, this is because of the time it takes to prepare and enshrine new knowledge in an educational context and, in some cases, because of the financial outlay involved. This is particularly the case in digital technologies concerned with prototyping and animation using the most advanced packages. In education at the university level, there were no courses specifically developed to cater for this requirement. In 1997, a wholly new studio was conceived to fill this gap at the Glasgow School of Art, one of the UK's most well known art and design institutions and home to the Art Nouveau architecture of Charles Rennie Mackintosh. The Digital Design Studio, as it is now called, had its genesis in the design faculty, but rapidly developed its own ethos and mode of operation in a site remote from the main campus.

The idea behind the studio was to create a leading-edge facility within the higher education sector in the UK, which would act as a flagship for 3D digital technologies. It was recognised, however, that the huge expense involved in creating such a facility would require a wholly different approach than normally employed. Thus, a strategy was developed at the outset to engage companies that were themselves at the leading edge of these technologies. SGI and Alias|Wavefront formed significant educational partnerships and a shared view of the overall aims of the studio. At the same time, industrial partners were sought, including The Ford Motor Company, the BBC, Unilever, Red Lemon, Picardy TV, and DERA (the UK Defence Evaluation Research Agency).

Because of the strategic role that the Digital Design Studio would have at a local, national, and international level, an ethos was developed which spread the studio's remit across three important and interlocked areas: education, research, and training.

EDUCATION

At the heart of the studio's educational activity lies the two-year multidisciplinary masters programme that was validated in 1997. From the outset, it was considered essential that this course would bring together graduates from a variety of backgrounds and disciplines. This strategy appears to have paid off, and to date graduates with the following backgrounds have been recruited to the masters programme: psychology, marine engineering, astrophysics, fine art, product design, marketing, animation (conventional), graphic design, and clinical science.

The broad aim of this two-year programme is to bring participants to a high level of capability in advanced 3D software, together with a critical sense of how to deploy it, particularly with reference to the subject

matter of their first degree. That is, if a student applies with a degree in engineering, we would expect to see the new skills developed focused on animation related to engineering subjects. The course was also designed as a platform for identifying and promoting PhD candidates, although it is worth noting that the studio has now recruited its first student already in possession of a doctorate. In recognition of the timeliness and importance of this course, a number of places are funded by external organisations, most notably the BBC and the British Defence Evaluation Research Agency.

The masters programme does not require prior computer experience as a prerequisite. As a rule, technology and all reference to it is subordinated to an "ideas-first" principle in which the technology used should be as transparent as possible. Year One of the programme is about learning new skills, and through imaginative staff-led projects, students acquire the basics of 3D modeling, shading, and animation. As expertise is accumulated over the course of Year One, students are asked to develop an individual programme of study to be undertaken in Year Two. This study, in effect a small-scale research programme, requires elements of the students' first degree expertise to be exploited through the newly acquired skills in 3D visualisation and animation. It is astonishing how rapidly people can learn and master what is often complex and convoluted software. The multidisciplinary nature of the course allows a cross fertilisation of learning styles and working methodologies. The staff within the studio are multidisciplinary in nature and are all originally from industry, with expertise in graphics, animation, product design, Web design, film making, computer game design, and technical programming and systems support.

To support students in devising and creating such a diverse range of outcomes, a number of educational and technical strategies have been developed. To ensure the necessary level of detachment and objectivity, the programme includes a research methods programme in which students are required to explore in detail the subject matter of their proposal and acknowledge the work which has preceded them in this area. Secondly, they must describe, through aims and objectives, the intention behind their proposal, taking into account the market for which it is intended. Formulating an appropriate technical strategy is vital to avoid needlessly carrying out unproductive work. Classes in storyboarding and critical precedents from film and TV are offered, the latter to provide a broad operational context.

It is recognised that teaching advanced software is time consuming and demanding, and that it has the potential to confuse. To overcome these problems, software workshops are followed with discussion "clinics" where students can informally seek guidance and support from various staff members. Of great importance, however, is the ability of the students to inform one another and introduce different patterns of thinking when engaged in project work. This is specifically encouraged through the medium of a student-led 3D forum where they discuss and show the work they have created.

RESEARCH

Physically adjacent to the masters studios are the research environments where the research focus represents the leading edge of digital 3D activity and is conducted in collaboration with several major companies. The benefits to students and their own individual programmes of study are immense, with many individuals or small teams being involved where appropriate.

Two main areas form the core of research activity:

1. Leading-edge product integration of new 3D digital technologies that support fundamental levels of the human computer interface. This research involves real-time advanced 3D visualisation techniques supported by gesture-based control, haptic feedback, and 3D audio rendering cues.
2. Development of virtual digital prototyping based on researched 3D reconstruction of objects, scenarios, and environments, permitting subsequent interrogation and interpretation of data.

Examples of research undertaken or still in progress range from virtually prototyping and animating early gliders in the quest for powered flight to development of new 3D workstations for automotive design.

TRAINING

In addition to disseminating advanced technology to industry, training activity serves as a vehicle for obtaining feedback into research programmes on areas of new design methodologies and human-computer interaction. The proximity of industrial users to the masters students permits a merging of learning and organisational issues, an important element in the overall development of the studio's ethos. Similarly, leading-edge research undertaken in the studio directly supports special training activities.

The essence of the above mix of education, research and training is to ensure the maximum cross fertilisation and exposure between masters students in the education programme, industry in the training programmes, and research activity undertaken for major UK and international companies.

THE DIGITAL DESIGN STUDIO ENVIRONMENT

To energise all of the above consistent with the high level of achievement anticipated, an exceptionally well-equipped studio had to be created. Paradoxically, the studio's location was to be another Charles Rennie Mackintosh called The House for an Art Lover. This building, although designed in 1900, was not built until 1995, and the Digital Design Studio was housed in the upper floors of this magnificent Art Nouveau building.

The studio currently has 40 SGI workstations ranging from Windows NT to O2s, Octanes, Origin 200, and ONYX IRs. Students have guaranteed access to a machine at all times. Other highly specialised items of equipment include the world's largest virtual reality table; a Fakespace Immersive Workbench; a Visionmaker digital draughting tablet; a Lake 3D audio system; a Monkey 2 skeletal animation system; haptic devices comprising gestural, touch and force feedback; and a Microscribe 3D digitiser. The studio also has a sound suite running Pro Tools and Logic Audio networked to video editing facilities. The nature of the agreements with SGI and Alias|Wavefront ensures that the studio is maintained at the leading edge in both hardware and software provision.

Much thought went into the layout of the studio to make the interplay of different facets visible, whilst at the same time acknowledging discreet activity where appropriate. The obvious consequences of research activity required a high degree of confidentiality, and the main research studio has access by security door entry systems. Confidentiality and non-disclosure agreements became a part of everyone's life in the studio (a reality of the world we live in).

From the foregoing, it will be seen that the studio has been designed to ensure the maximum physical and intellectual interplay among each of the three areas and to provide a highly charged learning and research environment. The operation of the studio does not follow the typical academic model. Its flexible hours and emphasis on the critical importance of deadlines more closely resembles the industrial model.

CASE STUDIES

In the second year of the masters programme, students engage in their own individual small-scale research programmes. This requires elements of the students' first-degree expertise to be exploited through their newly acquired skills in 3D visualisation and animation. The following four examples are typical of the diversity and type of subject investigated. It should be said that in each of the cases cited below, the student had no prior experience of 3D software. The work was created using Alias|Wavefront Power Animator and Maya.

Martin Usborne (psychology graduate)

Investigating the Nature of Digital Space

This submission was essentially an argument for a new perception of space through digital technologies. Usborne looked at precedents from film, theatre, and the painting tradition, and in particular at examples in which space is treated differently. The work of the German Expressionists, the theatre of Brecht, and paintings of Magritte formed part of his survey. This resulted in an exploration of digital space through two set pieces:

1. *Abstract Spaces*

Spaces that have only a vague sense of existence explored through light, movement, and shaders.

2. Momentary Architecture

Spaces that exist within a moment frozen in time. Porcelain plates are suspended from the ceiling of a cold, hard-floored room. As the plates drop to a shattering conclusion, the animation is stopped, and the spaces are explored (see Figure 1).

Brian Robinson (mechanical engineering graduate)

Visualisation for the Offshore Industry

Robinson started by looking at the current use of advanced visualisation technology in the offshore and marine industries and to suggest and demonstrate future developments in this field. He concluded from his research into the current use of digital animation that the most credible use of visualisation software in the marine and offshore industries was in training and simulation, and in particular, subsea visualisation. Based on construction drawings supplied by the manufacturer, a subsea plough used to lay transatlantic cables was built, and the sequence showing its descent to the ocean floor animated.

Alex Horton (product design graduate)

Avatar Choreographer

The object of this project was to produce a set of avatars that were correctly animated, optimised for real-time, and viewable through a VRML browser. Horton looked at emerging 3D shared spaces on the Internet and considered these in relation to the pioneering principles laid down by computer games. Horton used dance as an appropriate medium for animation because of the potential for social interaction and as a test of character design by virtue of the poses it offered with the normal walk, run, wave set of gestures. Break dance, folk dance, and rock n' roll were studied. The outcome was the creation of VRML avatars capable of performing a variety of dance moves (see Figure 2).

Tom Marshall Andrews

(product design engineering graduate)

Bus Transportation Interface, in conjunction with DERA

Although it is part of the UK Ministry of Defence, DERA has promoted new products, materials, processes, and advanced technical development leading to many major technologies used worldwide in everyday products. Based on its significant scientific, technical, and engineering experience, DERA recognised the need to build rapid and convincing digital prototypes which could be seen operating in context, albeit virtual. Working with DERA, Andrews's concern was to prototype product concepts based on a new type of display material with unique operational characteristics. Once the concepts were prototyped, the key to further development was to concentrate on the whole user experience and promote the product in a "real-life" environment. This was achieved by shooting video of a real streetscape complete with bus stop, traffic, passers-by, and a typical user, who would go through the motions of interacting with the new product.

After the prototype was modeled in 3D, it was carefully composited into the scene. The bulk of this digital mock-up still had to have the very necessary visual cues that would, in essence, bring the product to life. These subtle, but important, cues simulate the sensory cues that we experience in our everyday lives.

The time required to produce this type of high-level presentation is short, which gives fast concept studies, market surveys, or customer focus groups a real benchmark in terms of feedback and interpretation. Yet many organisations choose not to use this approach. It is important to realise that nothing was produced physically. This must beg the question: Do you present virtual products and only produce what you have orders for?

NEW RESEARCH AREAS

The small-scale research projects undertaken by multi-disciplinary masters students have contributed to the overall research profile of the studio. As evidence of this, the studio has recently been awarded a £1.1 million (\$US 1.76 million) research grant from the Engineering and Physical Sciences Research Council to look into the use of advanced 3D displays in selling automobiles. This is the first time that research projects have been funded by this research council in an art and design institution in the UK. This in itself is an example of the growing interest in the sensibility that such institutions can bring to "hard" research issues. This research, entitled Supporting Customer Interaction, Feedback and Information (SCIFI), seeks to develop interactive 3D displays and sound systems, which will allow the public to engage in pre-sales interaction with realistic and credible computer-generated models, resulting in a fundamental reassessment of the way in which cars are sold. As such it is the culmination, to date, of the distinctly multidisciplinary ethos promoted within the studio.

In carrying out this research, the studio will be concerned with developing future interactive selling strategies in which advanced 3D computer modeling and virtual environments with touch and sound cues, together with communication technologies, will allow new and imaginative sales scenarios to emerge.

The present method of selling cars requires investment in a physical showroom, which nevertheless carries limited stock. In addition, the showroom environment can be pressurized and sometimes driven by what the sales person needs to sell rather than what the buyer wants. The challenge presented by this research is to present complex information for easy public access. This can be achieved through a non-technical, non-intimidatory approach, thus encouraging a greater degree of social inclusion.

It is hoped this research will identify appropriate visualisation technologies, which will fundamentally improve communication processes between customer and manufacturer. In turn, this will result in the public having access to information of a qualitative nature delivered in a visually exciting and credible manner. This will enable the manufacturer to move from the present showroom convention to having a "stock list" presence in densely populated retail environments, such as airports, stations, or shopping malls.

A Multidisciplinary Approach to Masters Education in Digital Research

CONCLUSION

The multidisciplinary approach towards research taken in the Digital Design Studio has, over the past two years, produced some remarkable results that could not have been predicted under the constraints of normal single-disciplinary thinking. It has also contributed significantly to winning funding from a number of industries as well as a major UK research council.

In general terms, it is well established that a group of people from different backgrounds can together act as a powerful lever in imaginatively lifting ideas off the ground. Such a resource provides a powerful think tank, which can solve many problems and generally create a situation in which many approaches and conclusions are possible.

In the context of this paper, the role that multi-disciplinary education can play in new digital industries is recognised. Game companies, Web development organisations, digital online services, film and TV, and manufacturing and visualisation industries demand these empowering combinations. It is the experience of the Digital Design Studio that putting a psychologist together with an engineer, a designer, and a fine artist can lead to wildly exciting, if sometimes unpredictable, outcomes.




Figure 1: Martin Usborne's exploration of digital space using porcelain plates.



Figure 2: Alex Horton's dancing avatars.

WHAT IS NETPAINTING?

NetPainting describes the use of a feature, NetPainter, in the paint software Painter that supports real-time (synchronous) collaboration between Painter users across a local area network or via the Internet. During a NetPainting session, as one user paints or makes changes to an image the other users see the brush strokes and effects appearing in the canvas window on their own computer screens. Only one user can paint at one time. When users finish painting, they turn the controls over to another user by means of a "traffic light" system. Another user can then carry on where the previous user left off. Users take turns working until the project is finished and/or the session is over.

A neat feature of NetPainter is the Chat Window. All users, at any time in a NetPainting session, can write short notes that appear in every user's Chat Window. In this way, interactions can be written as well as visual.

HOW THIS COLLABORATION CAME ABOUT

This collaboration was prompted by Jeffrey Kraus, an art teacher at Anderson High School in Cincinnati, Ohio. He came across my Web site (www.jeremysutton.com) and sent me a note asking about my interest in joining a NetPainting session with some of his students. I was enthusiastic about the idea. I was keen to explore the artistic interaction between myself and the students in a way that would be educational as well as fun. The students were also very excited.

Jeff's objectives were to grow his students' creativity, enrich their visual awareness, and expand their horizons by exposing them to outside professional artists in the working world.

THE EQUIPMENT

I used a Macintosh 9500/132 (228 MB RAM, 2 GB drive) with 20-inch color monitor, 28 baud Supra modem, and regular phone line. The school used Pentium PCs (64 MB RAM, 2 GB drives) with 15-inch color monitors and T1 lines. All computers had Painter 5.03 and Wacom tablets installed. We worked on canvases that were 8 inches x 10 inches at 72 DPI.



From Session 1 (self portraits)

SESSION FORMAT

Jeff and I agreed on a tentative format for the sessions. The initial format was three 30-minute sessions, each with three students and myself.

Session 1 would be loosely based on self-portraits. Each artist would establish an overlapping rectangular area within the compositional area (each rectangle should overlap the others about a third). Each in turn would develop a self-portrait (using a mirror placed next to the computer), making sure to address the composite areas.

The second session would be a single portrait that all artists would work on, using the following suggested steps.

1. Establish a textured or patterned background.
2. Use distort or similar brush to lay out image.
3. Establish value pattern.
4. Begin building details.
5. Develop a decorative border.

The third session would be freeform, beginning with expressive brush strokes and adding on to the composition either by moving towards imagery or maintaining a non-objective abstract composition.

These general guidelines were given to the students:

- Work for about three-four minutes at a time.
- Use the brushes. Stay away from special-effects filters.
- Add on to or modify an image but don't clear the canvas and eliminate all previous work.

THE EXPERIENCE

We set up a technical rehearsal session followed a day later by the "real" thing. After a few initial tries, we succeeded in establishing an Internet connection between my computer and three computers at the school.



From Session 2 (single portrait)

Painting Together Across the World: Collaborative NetPainting Between High School Students & Professional Artist

I initiated the painting in most sessions and then passed the control on after a few minutes. It was interesting to communicate with the students through the Chat Window. We encouraged each other, commented on the brush strokes and tools and effects we saw being applied, discussed the experience, and talked about creating art in general.

Jeff describes the experience at the high school:

“This experience went far beyond my expectations. I felt that the process was artistically stimulating, and the results were outstanding. I wish you could have been here to have seen my students as they worked. They were engrossed in the process, and each and every one of them were trying hard to watch what you were doing and how you were doing it. The only major problem that I had was too many students wanted to work with you (Jeremy there obviously is not enough of you to go around).”

THE RESULTS

You can see here results from the three sessions. I encouraged the students to save versions regularly and thus build up a documentation of the process.

THE BENEFITS

There were clear artistic, technical, and social benefits for both myself and all the students involved. In this section, I let Jeff and his students speak for themselves.

ARTISTIC BENEFITS

Jeff

“Creatively the students were forced to become more flexible because they had to respond to changes to their work that was beyond their control. The quick pace of the process forced them to become more direct and react more intuitively.”

Christy

“I generally work tight, and I learned to be more free when painting. It has also helped me to be less conscious when people are watching me work.”

Eamon

“I learned that as an individual my ideas became narrow, but when I was working collaboratively my creative thinking was greatly enhanced.”

Sarah

“I liked how the style and techniques of four people all came into one picture.”

TECHNICAL BENEFITS

Jeff

“Technically, the students learned practical techniques: how to use specific brushes, how to construct an image from the very general to the specific, how to employ subjective color and still maintain a sense of reality.”

Sarah

“I enjoyed watching how other people use the tools provided, and I believe I understand the use of creation tools better.”

Whitney

“Some techniques and skills I acquired were shading for facial features.”

Eva

“I learned how to use some new brushes, how to build up a face using different media.”

SOCIAL BENEFITS

Eva

“I learned how to work with others.”

Whitney

“I learned that working with others is fun, but not as easy as thought to be. A lot of cooperation and patience is needed.”

Rory

“I learned the importance of collaboration in artwork. Waiting to take a turn makes one realize the importance of everything contributing to one whole. It was great seeing how the work evolved as a whole.”

Whitney

“I think the fact of more than one person owning the painting and many people contributing to one piece of artwork is astounding. This way, no one person can be judged for what they did or didn't do.”

Eamon

“I love the idea of working with someone from across the country. Hopefully someday I will be able to work across the globe.”

Lindsey

“I liked putting a piece together with someone that could be anywhere in the world. This enables people from anywhere to present new ideas and styles to you.”

Rory

“I liked working with someone other than my peers and someone who had artistic experience. It was a great learning experience.”

Christy

“It was really neat to work with a professional and get feedback.”

Whitney

“Working with Jeremy Sutton on Painter was so wonderful. I considered it an honor to work with Jeremy Sutton because creating a piece of artwork with an artist as admirable as him was a phenomenal experience.”

THE CHALLENGES

Logistical

We were three time zones apart and thus had to carefully coordinate our session schedules so we were ready at our computers at the same time. On the student end, this was an extracurricular art department activity. Students gave up personal time to participate.

Technical

Initially, there was a fire wall problem at the high school that had to be solved. We then had to carefully follow the set up procedures for connecting through Painter. We also had to master the controls and techniques associated with using the NetPainter feature, such as the traffic-light buttons. Sometimes the connection was lost inexplicably and the session abruptly ended! In those cases we reconnected and started again. I found it useful when students customized the User Name and Chat ID in the Set Up dialog box. This allowed me to easily see which student was painting and who was sending chat messages.

We were careful to ensure that all the Brushes and Art Materials in Painter on all the collaborating computers were the same. Even so, there were still some problems:

Whitney

"It got confusing with certain brushes and brush strokes that didn't show up on certain computers."

Structural and Temporal

How to structure the sessions? How many students to work with in a session? How long should a session be? How long to allow each user to work before passing on control? How to ensure that everyone took their turn in the right order? All the students wished there were fewer people and more time per session. Several students said they would have preferred to work one on one:

Eva

"I only got about three minutes to work, and I was a little confused about how long I was supposed to take and whose turn it was. I think it would be more efficient to work on a one-to-one basis, so that we could get more done, and it would be less confusing."

In subsequent sessions we listened to the students' feedback, reduced the number of users in a session and increased the duration of each session.

What constraints, rules and boundary conditions to impose on the process and content? The constraints we adopted were designed to challenge the students and exercise their artistic skills.

THE FUTURE

We are looking forward to continuing the NetPainting collaborations, and Jeff and his students are hoping to link up with other schools and artists around the world (e-mail Jeff at jkraus@foresthills.edu if you are interested).

I hope the success of my NetPainting collaboration with the Cincinnati students serves to encourage schools and artists to form their own NetPainting partnerships and thus help develop student's artistic, technical, and social skills.

Jeff says it best:

"I wanted to express to you what an impact you had on my students. Each and every one of them felt they had grown as a result of their contact with you. Professionals do not realize what impact they have on students. It is professionals like yourself taking the time out for today's youth that makes what we teach in the classroom concrete."

Besides expanding and stimulating them creatively, the NetPainting experience also leads the way naturally into the reality of 21st-century telecommuting. Many of these students will work on computers at locations separate from their employers.

Here is a real-life opportunity for students to interact with, and learn from, someone they may otherwise never meet. They will take that example as a model for the way creative collaboration can be achieved via the Internet between artists anywhere in the world. NetPainting is clearly a bridge to future realities.



From Session 3 (freeform)

INTRODUCTION

Many students entering college who choose to major in computer science do so because of their exposure to graphical environments (including games). Although it is not always the primary focus of a graphics course, development of a simple game or similar project prepares students for developing and producing graphically oriented products. While the focus of the graphics course presented in this paper is one of traditional graphics topics, user interface principles, and the extension of background topics (i.e., data structures, algorithms, architecture, etc.) to graphical environments, a major term project is required. The primary emphasis of the project is on developing a small product rather than just completing a class exercise.

The next section discusses the course outline and an overview of the project. The following section presents information on activities required to support such a project. Next, the project framework is presented along with observations. Finally, concerns for future enhancements are addressed.

COURSE OUTLINE

The computer graphics course presented here is a senior-level elective with traditional background requirements including data structures, algorithm analysis, and mathematics (trigonometry and matrix algebra are recommended). The course format has been fairly traditional, with topics including an introduction, hardware, software, rendering, modeling, theory, and advanced topics as time permits. The current course format is the result of an evolution from a traditional 2D-based course with some 3D at the end to a 3D-centric course.

The current 3D approach to the course grew out of experiences of the SIGGRAPH 97 conference,¹ where Edward Angel² presented a course model that indicated students could achieve 3D rather quickly. Since 3D had been difficult to achieve in previous course offerings, his outline was adopted, and in particular his suggested series of exercises³ utilizing the “maze” theme (three assignments) was implemented in the 1997 course with success. Students were able to accomplish 3D projects by midterm.

As in most successful situations, it was noted that the difficulties “moved.” Students now struggled with lighting, texturing, and more advanced topics. An enhancement for the 1999 course included an extension of the “maze” series to include five assignments in addition to the project. The extensions were a lighting/shading lab and a texture lab. Additional activities (homework) were assigned to introduce OpenGL, practice modeling techniques, and provide theory-based problems. A major goal of the course was to complete a project. The timeframe for the project was approximately one-half of the semester, although project requirements were given on the first day of the course. A project proposal was due at midterm, an intermediate report was required at approximately three-quarters of the way through the semester, and the final report was due before the end of the term.

PROJECT SUPPORT

The laboratory assignments were designed to develop facility in the use of OpenGL, explore reasoning and representation using geometric objects, and to provide a small “engine” that might be useful for the project. Each assignment depended on completion of the previous assignment through the first three assignments. Each student assignment was captured on videotape and given to the student for review. Some of these were played in class and evaluated by all of the students.

The first assignment was to draw a 2D maze on the screen and recursively remove the walls to produce a connected maze.³ A hedge was required to encompass the maze, and although the entrance side was chosen at random, the exit was forced to the opposite side. This exercise provided some OpenGL programming experience, including simple primitive and attribute manipulation. It also introduced students to an algorithmic technique that would be useful later. In addition, careful reasoning about special situations (boundaries and adjacencies) provided the opportunity to develop and extend reasoning skills. Much of the difficulty associated with completing the lab involved these reasoning problems and helped reinforce the idea that core concepts developed in previous courses were essential even in “graphics programming.” Most students completed this assignment successfully, although some did not complete the recursive removal portion and had to “back-patch” a temporary solution.

The second assignment was to add a “rat” to the maze³ and then manually traverse the maze. This involved positioning and movement of an object through the maze. Graphically, this exercise provided students with the opportunity to erase and redraw (back buffer) objects and to develop a simple collision-checking algorithm (again, somewhat more traditional programming). In addition, working with interactive events utilizing the keyboard and mouse in a window provided an introduction to interface issues (for example, no middle mouse button, etc.). Many of the students were successful in completing this exercise, although the heaviest class attrition occurred at this stage.

The third assignment was to develop a 3D “rat” view of the maze.³ This changed the 2D lines to 3D surfaces and introduced a perspective view. Motion was achieved by changing views (no smooth motion). In addition, the original 2D window was retained as a map of the maze, requiring the students to control multiple windows. Again, most of the students were successful, although there were some interesting problems (for example, left and right mix-up, slow operation due to swapping buffers too often, etc.). A sample display is shown in Figure 1.




Figure 1. 2D map/3D view.

The fourth assignment was to add lighting and at least one shadow to the 3D portion of the maze. Most students tried to accomplish this through material properties, which generated “soft” shadows. Some students created an object (for example, cheese at the end of the maze) to produce a “hard” shadow, and the remainder created “hard” shadows using the walls of the maze. An interesting visual effect was produced when some students produced a darker color at the top of the wall than at the bottom.

The final assignment was to add texture to the 3D maze. Several students used the sample checkerboard pattern⁴ presented in class for their texture. One student developed a utility to capture textures from files and place them in the program as a data object. A sample “final” maze view is shown in figure 2.




Figure 2. Shadow and texture.

THE PROJECT

The project encompassed the last half of the semester, although it was assigned on the first day so that students were aware of the project and could begin thinking about what they might do. Also, it was suggested that students work in teams to give them the opportunity to discuss common interests. It would not be fair, however, to consider that the project covered the entire term since until laboratory assignment three was completed, many students did not believe that they could work in 3D.

The project assignment was divided into three parts: a proposal, an intermediate report with a working prototype, and a final report with the working project. In support of the project, three information resources were utilized. The first was an invited speaker from the art department, who presented some design fundamentals and illustrated them with examples. Most of the students found this information helpful in their project design. The second was a presentation from a marketing department professor on product design, development, and related issues. Although it was more difficult to incorporate, students found this presentation interesting and, as a class, requested a copy of all of the projects (which was provided with labels, executables,

source code, and documentation on a CD-ROM at the end of the term). The third information resource was a set of handouts from written and online references concerning story telling and writing. The success of this material was mixed, since it was given late in the course timetable and without a class guest presentation. Each of these topics was introduced to explore and support a non-programming aspect of the project.

The proposal covered the project outline, a general design, an implementation timetable and sample screens or views. In many ways, this was the most challenging part of the project. Several students encountered difficulty in choosing a topic, and in general, all had trouble with the scope of the project (it was either too limited because they were not confident about working in 3D, or it was too ambitious because of the commercial products they were accustomed to using). Only a few students had any working code at proposal time. Sample projects included games, walkthroughs, and stories. All of these projects utilized techniques and code segments developed in the lab exercises.

The intermediate report included an overview, a preliminary user guide, an issues section discussing design or implementation problems that might limit the scope of the project, an updated timetable, user testing, and a working prototype. This report was critical to the final success of the project. In particular, if the project was off track or too incomplete, feedback from this report provided the student with some guidance for correction and limited refinement of the project goals.

The final report was a small package including documentation and media. The report included an introduction, a complete user guide, a project summary including time spent, algorithms utilized, limitations and unaccomplished objectives, references, and appendices with listings, artwork, and support files. The media contained online documentation, source and executable code, and a soft copy of the report.

OBSERVATIONS

The course was successful in introducing students to “graphics programming” and a project-product approach. Based on faculty analysis and student feedback, future course offerings will need to increase the emphasis on modeling, provide additional support for non-technical issues (while retaining both the theoretical and implementation content currently in the course), incorporate an object approach, and encompass additional advanced techniques.

References

1. SIGGRAPH Conference Grants for Educators Program, 1997.
2. Angel, E. (1997). Teaching a three-dimensional computer graphics class using OpenGL. *Computer Graphics*, 31 (3), 54-55.
3. Angel, E. (2000). *Interactive computer graphics* (2nd ed.). New York: Addison Wesley.
4. Woo, M., Neider, J., Davis, T., Shreiner, D. (1999). *OpenGL programming guide* (3rd ed.). New York: Addison Wesley.

One of the key points in creating an effective computer science course is developing the projects students will do in the course. This paper describes a set of projects for a computer graphics programming course, with a particular focus on projects whose content is in the area of visualization in the sciences.¹ This focus serves a well-defined group of students whose needs have long been a concern to the SIGGRAPH community² and can also make the course a key factor in a computational science program. The projects are intended to allow the course to be both an effective service course for science students and a sound introduction to computer graphics for computer science students.

THE COURSE

As described elsewhere,^{2,3} the author believes that developments in graphics hardware and APIs allow creation of a new paradigm for the introductory computer graphics course. This course can now focus on graphical thinking, problem-solving, and communication as expressed through a high-level API, and it can provide an excellent background to students who would want to use computer graphics in their professional work. This new course is a half-year course that serves as both a service course for many other programs on a university campus and as an entry point for students who want to study computer graphics more deeply. Its only requirements are programming skills (as expressed by successful completion of the equivalent of the CS1 and CS2 sequence) and an ability to think geometrically and spatially in the student's field.

One of the key issues in creating any successful computing course, and one that often is overlooked in the literature on teaching computer graphics, is the projects students will do in the course. This is especially true of a course that offers significant content to non-computer-science majors, because the projects are the students' primary tool to encounter the subject matter.

PROJECT CONCEPTS FOR THE COURSE

The author suggests that a half-year course would include five programming projects. These projects are intended to develop from simple geometric thinking to more significant applications of the tools in the API. Note that these projects do not identify the API being used; it should be possible to do them with any of the current APIs that might be used in the course. To begin, let us outline the projects in very general terms, focusing on the topics to be considered without specific applications (these will follow in the subsequent section on science applications).

Project 1: Identify some appropriate geometry to be displayed using the API that supports the course. This requires proper initialization of the API, definition of the viewing environment, analysis of the geometry into appropriate primitives and use of simple instancing transformations, use of geometric primitives to display the geometric primitives, use of hidden-surface display, and use of color as needed for the geometry. A technique such as the OpenGL idle callback can be used to create some motion, even in this first project.

Project 2: Use the same or different geometry, but increase the sophistication of the rendering by adding shading and lights to the scene, illustrating ambient, diffuse, and specular lighting. Use multiple lights, colored lights, or spotlights to illustrate the effects they can create.

Project 3: Use the same or different geometry, but focus on simple interaction with the scene, adding keyboard, menu, and mouse callbacks and controls for the scene. If tools are available (such as the MUI system for OpenGL) add button or slider controls as appropriate. The approach to this project will depend on the amount of UI focus the instructor wants to have in the course, but this could be significant.

Project 4: Use the same or different geometry and add more features of modeling and/or rendering, including hierarchical modeling, clipping planes, alpha blending, fog, or texture mapping. This is a good place to discuss genuine 3D viewing with stereo pairs or Chromadepth color encoding,⁶ and this is very interesting to students. While it is tempting to use techniques for their own sake, this might be a good project to leave fairly open-ended to allow the student to choose (and justify) the techniques to be used to create a particular communication.

Project 5: A more advanced, "capstone" project would ask the student to use additional techniques noted in Project 4 that were not used there, as well as new techniques, such as evaluators to generate spline surfaces or object selection to allow students to expand their interactions with images. Again, this project could be built on the same or different geometry as earlier projects, as appropriate.

SOME SPECIFIC PROJECTS

The programming approach to the computer graphics course with science-focused projects offers an excellent opportunity to create a course that has unique value to students outside computer science. In particular, focusing on science topics can make this course effective for science students or, especially, computational science programs. This section describes some such projects and gives examples of the kinds of images each could create. Restricted space limits the number of projects that can be described here, so we describe only projects for mathematics, chemistry, and physics. All of these projects, as well as others, are described in more depth in the online materials noted at the end of this paper.

Project 1

- **Mathematics.** Display a surface defined by a function of two variables. This requires creation of a grid in the domain and corresponding points in 3-space in the surface, displaying the triangles derived from the domain grid, and coloring the surface in some way such as height.
- **Chemistry.** Read a standard molecule description file and create a display of the molecule with standard atom-coded colors. This requires the use of simple transformations to place the atoms, using pre-written functions to read molecular description files.
- **Physics.** Display a rectangular bar divided into a grid, with some grid areas held at constant hot or cold temperatures and with temperatures changing by a diffusion process. Because temperature can be encoded through either geometry or color, this project offers the opportunity to discuss ways to encode information in a graphics environment.

Project 2

- **Mathematics.** Add shading and lights to the surface viewing. This project could also focus on higher-dimensional or parametric surfaces, generalizing the first project.
- **Chemistry.** Add shading and lights to the molecular image, showing the atoms with appropriate highlights.
- **Physics.** Display a function defined as the sum of two linear or circular wave functions, with ambient, diffuse, and specular lighting. Use the simple animation of the first project to show the behavior of the wave functions over time.

Project 3

- **Mathematics.** Add keyboard control for rotations, clipping, or other mathematical operations. This would include discussions of the kinds of interaction or control appropriate for mathematical operations. Add other kinds of control for parameters if the surface is defined as a parametric function.
- **Chemistry.** Add keyboard-controlled rotations and menu selections to the molecular display. This would allow discussions of the kind of control appropriate for viewing objects in space. Because of the time it might take to display fairly complex molecules, the third project should also add display lists to improve performance on the display.
- **Physics or Chemistry.** Create a simulation of molecules moving randomly in a cube, and periodically report the number of times a molecule strikes the side of the cube in a step of the simulation. This value simulates pressure, and allows the student to test the gas laws: the pressure varies linearly with the number of molecules present and inversely with the volume of the cube.




Figure 1: (a) A mathematical surface with light and materials modeling (Project 2).



(b) A stereo pair of conic sections that the reader can resolve without external aid. (Project 4).

Have the student track the motion of one molecule to see the nature of Brownian motion, allow rotations of the space so the user can see the entire space, and add controls that change the number of molecules present or the volume of the space to allow the user to test the simulation.

Project 4

- **Mathematics.** Create a display showing the quadratic curves as actual conic sections, using menus and other controls for managing the clipping planes that create the sections and creating coordinated stereo pairs to see the conic sections in 3D.
- **Chemistry.** Add transparency and alternative views of the molecule. The alternative views could add a clipping plane with keyboard control of a front-and-back motion on the plane, Chromadepth texture mapping to create real 3D viewing, stereo pairs, or texture-map information onto individual atoms, though this may be very slow on some machines.
- **Physics.** Create a display of a function of three variables with a physical basis, such as temperature, static force, or gravity in some spatial region. This could be an implicit surface or a color-coded display of function values in a clipping plane in the region. Allow the user to interact with the display to move the implicit surface or clipping plane, and have the project include some conclusions about the nature of the function.

Project 5

- **Mathematics.** Display a spline surface with user-selectable, movable control points. This requires students to manage both evaluators and object selection and to deal with moving objects in 3D.
- **Chemistry.** Add selection of an individual atom so that the atom can be manipulated or information on it can be returned. Alternately, the project could display vibration modes for a selected molecule component using vibration information provided separately.
- **Physics.** Create a simulation of a multiple-body problem with simple numerical integration and build a time-based animation of the situation. Use simple controls to allow the user to move around in the animation space, and use Chromadepth or stereo pairs to allow the user to experience the animation in real 3D space. Add selection to allow the user to identify an individual body and to use interaction techniques to provide real-time steering of the simulation by changing the velocity or direction of the moving body.

These figures illustrate some of the projects described above. The figures start with mathematics projects, then move to chemistry, and then to physics, and the captions identify which project is being illustrated. Note that some of the projects, especially Projects 3 and 5, do not lend themselves well to figures because their focus is an action, not an image.




Figure 2: (a) Simulation of the gas laws with sample text output from a query (Project 1).



(b) An image of a molecule with Chromadepth coloring added (Helvetane) (Project 4).

CONCLUSIONS

This approach to the beginning graphics course is oriented toward the science student and fits the growing use of visualization in the sciences, but it should be possible to focus the projects towards many other fields to support visualization in these fields. Versions of this course have been taught twice (at two different universities) with significant student interest and success, although it has so far been difficult to reach significant numbers of science students. Some stability in the course content and some marketing of the course to the sciences are needed to increase interest. Conversations about and presentations of this course to a number of people in computer science, computational science, and the sciences lead the author to believe that this approach is worth sharing with the graphics community in spite of the relative lack of experience, because other fields are eager to gain the advantages of learning about graphics without having to master fundamental graphics principles and without having to rely on pre-written graphics applications.

The project materials described here are being developed with funding from the National Science Foundation. These materials include descriptions of projects for each of the sciences, source code for the examples, supplements on various graphics topics, and pointers to example materials from various sources. As additional materials are developed, they will be available at www.cs.csustan.edu/~rsc/NSF/

and other pages linked from there, and the final results will eventually be available at the online SIGGRAPH and Eurographics sites and at appropriate courseware repositories.

References

1. Brown, J., Cunningham, S., & McGrath, M. (1990). Visualization in science and engineering education. In G. M. Nielson & B. Shriver (Eds.), *IEEE Tutorial: Scientific Visualization*. IEEE Computer Society.
2. Cunningham, S. (2000, February) Re-inventing the introductory computer graphics course: Providing tools for a wider audience. *Computers & Graphics*.
3. Cunningham, S. (2000). Powers of 10: The case for changing the first course in computer graphics. *Proceedings of the SIGCSE 2000 Technical Symposium on Computer Science Education*.
4. Reports of the Graphics & Visualization 99 Workshop are to appear in *Computer Graphics and Computer Graphics Forum* and are available online at www.eg.org/WorkingGroups/GVE/GVE99 and www.education.siggraph.org/conferences/GVE99/.
5. Domik, G. *Report on visualization curriculum*. www.uni-paderborn.de/fachbereich/AG/agdomik/visualisierung/vis-report/
6. Chromadepth glasses and other information are available from Chromatek Inc: www.chromatek.com/

Credits

This work is partially supported by National Science Foundation grant DUE-9950121. All opinions, findings, conclusions, and recommendations in this work are those of the author and do not necessarily reflect the views of the National Science Foundation. The author would also like to thank the San Diego Supercomputer Center for supporting this work.





Figure 3: (a) A heat simulation in a bar (Project 1).



(b) An N-body simulation (Project 5).

The QuickWorlds program, a collaboration between university researchers and elementary school teachers, provides elementary school instructors with virtual reality models that enhance the K-6 curriculum.

Almost 100 years ago, the Parent-Teacher Organization at Abraham Lincoln Elementary School in Oak Park, Illinois bought a stereopticon for use by the teachers in the belief that a stereoscopic display might give students insights not available through conventional imagery. Today, researchers from the University of Illinois at Chicago are collaborating with the staff of Lincoln Elementary to explore what learning benefits virtual reality technology may offer children within a school context.

Our long-term goal is to help prepare the nation's schools for the advanced visualization technologies, including VR, that are becoming operative in research laboratories. These technologies offer additional representational bandwidth and support new kinds of interactivity. Whether they will be adopted depends on whether we find more solid evidence of effectiveness for the individual learner, and on whether we can successfully align their use with school goals, curricula, practices, and culture.

In August 1999, we installed an ImmersaDesk in the media center at Lincoln, where it will remain for the next two or three years. An ImmersaDesk in the school allows us to continue our formal learning studies of how VR may benefit conceptual learning,^{1,2,3} work with entire classrooms on learning scientific inquiry skills,⁴ and investigate how this technology can be integrated into the existing school curriculum. This paper describes our QuickWorlds program, in which teachers request VR models to supplement regular curriculum. In this program, we serve as modelers; the teachers control the selection, specification, and use of the models.

Lincoln is a K-6 elementary school in Oak Park, Illinois, a racially and economically diverse inner-ring suburb bordering Chicago's West Side. Lincoln is attractive as a research site for its size, diversity, and state of technology adoption. It is a large school (620+ students), nearly always allocating four 20-to-30-student classrooms at each of the K-6 grade levels. In addition to a racially and economically diverse student body and faculty, Lincoln offers diversity of subject mastery, as reflected in state and local achievement tests administered at the school. While performing moderately above average as a school, Lincoln has significant representation in all performance quartiles. The school is also about average with respect to technology infusion, with about one computer for every five children, distributed both in classrooms and computer lab settings, and an orientation more toward computer literacy and technology education than conceptual learning.

VR IN THE SCHOOL

While the history of meaningful technology adoptions in schools has been discouraging,⁵ we have the advantage of the lessons learned from the computing integration efforts over the past two decades. The central theme of that experience has been recognizing the importance of addressing the needs of the individual learner while respecting the constraints imposed by the educational context within which learning occurs.

Among the factors that limit successful integration of advanced technologies in schools⁶ are:

- Insufficient resources to support teaching practices constrained by conventional school organization.
- Lack of alignment between technology-based materials and school curriculum/performance goals.
- Lack of authentically motivated teacher preparation and training offerings, and lack of time for teachers to pursue additional training.
- Insufficient technical support, both for maintaining and upgrading hardware and software systems, and for providing assistance on the operation and capabilities of application software.
- Failure to provide specific mechanisms for assessing the pace and effectiveness of technology integration with respect to student learning and school climate.

The uniqueness, cost, size, and fragility of VR aggravate these problems. VR technology will be deployed in considerably fewer numbers than personal computers. The existing learning application base is almost non-existent. There are no standard staff development curricula or off-the-shelf courseware. And technical support requires very specialized expertise.

Our goal is to align our project activities closely with Lincoln's learning goals, curriculum, and practices, so the teachers will feel confident that participation does not add irrelevant activities to an already crowded curriculum. Issues of fairness are also paramount. We plan to give every child in the school at least one VR experience per year, either in one of our larger studies or through the QuickWorlds.

We set up the ImmersaDesk in a room adjacent to the school's Media Center. This allows us to minimize the impact of "pull-outs" from regular classroom activity. Since classes regularly visit the school's Media Center, it is easy to "pull over" students to work individually or in small groups while the rest of their class is doing other activities in the Media Center. An SGI Deskside Onyx commonly drives the ImmersaDesk. Our software uses the CAVE library and SGI Performer. Since the CAVE library and SGI Performer are now available for Linux, we are also investigating the use of a Linux-based PC and passive stereo projection as a route to an eventual low-cost setup.

In preparation for the QuickWorlds program, we had a meeting with all of the teachers at the beginning of the school year to describe our goals, and to encourage their participation. The staff and teachers already knew us from our previous work with a temporary ImmersaDesk setup the year before. During the fall, we met with the local Parent-Teacher Organization and the school district's technology council to show them the ImmersaDesk, discuss what we were hoping to achieve, and answer any questions that they had. We also used this time to distribute permission slips to all of the parents, and to talk to interested teachers about developing several initial QuickWorlds. In the spring, we began regular hours at the school, with personnel on-site every Friday so any interested teacher could make use of the VR facilities.

QUICKWORLDS

The QuickWorlds program is intended to provide a fast-turnaround mechanism for teachers who would like to make virtual models available to their students as part of the regular learning program. Educational materials abound with visualizations for good reasons. Visual representations support learning in numerous ways and the more powerful the visualization technology, the more complete the support:

- *Integration.* A visual representation provides simultaneous and parallel access to multiple parts or components of the visualized system, while a text necessarily presents them in some sequence, leaving the learner with the task of integration.
- *Dynamics.* A dynamic visualization demonstrates how components of a system change over time and interact.
- *Reification.* A visualization can convert abstractions into perceivable objects and engage perception in support of conceptual learning.
- *Activity.* Interactive visualizations allow a learner to manipulate a system, and they draw upon the near-universal principle that knowledge is constructed in the course of activity.
- *Immersion.* High-end visualizations combine all of these features to let learners feel as if they are directly experiencing the visualized system, thus drawing upon children's natural capacity for experiential learning.

However, recognizing the power of advanced visualizations to support learning does specify which visualizations are most useful for teaching a particular topic. Working closely with the teachers and creating the visualizations that they believe will be beneficial helps direct this inquiry. During regular meetings, the teachers submit requests for models, specifying the object(s) to be modeled and the facets of those models most important to the underlying learning goals. One of our art students then creates the necessary models, and our personnel assist the teacher in the presentation of the models to his or her students.

The first set of QuickWorlds that we have developed covers a wide spectrum. Each of these worlds took one to two weeks to develop, with the bulk of the development effort focusing on building the models. The models are then loaded into a common viewing program that allows stereo head-tracked viewing and interaction. This common viewing program allows several sites to remotely collaborate in viewing and manipulating the same model — a capability that we plan to leverage in future years. The first six QuickWorlds are shown in Figure 1.

The models that have been developed range from simple static models such as a wood ant and the interior of the Earth, to more complicated dynamic models such as the volcano, iceberg, human heart, and solar system.

- The ant was created as part of the school's units on insects. It is magnified to be six feet long, to make the 3D structure of the ant easier to see.
- The interior of the Earth, created on request of a fourth-grade teacher, shows the Earth as a four-foot-diameter sphere. Kids can use a dynamic cutting plane to cut through the planet and see the internal structure.
- The volcano, created as part of the school's unit on earth science, shows Mount St. Helens before and after the eruption of 1980. It allows the children to look under the surface of the ground, and to use the cutting plane to see what the internal structure of the volcano looks like.
- The iceberg was created as part of the school's unit on the polar regions, allowing the students to see the structure of a small iceberg both above and below the water. Animated seals, penguins, and whales help give a scale to the iceberg and lead the children from the visible iceberg to the larger mass of ice beneath the water.




Figure 1: A selection of QuickWorlds that have been developed. Top row: a wood ant, the interior of the Earth, the human heart. Bottom row: an iceberg, the solar system and a volcano.



Figure 2: Kathy Madura, physical education teacher, discussing the human heart with her students using the ImmersaDesk at Lincoln.

QuickWorlds: Teacher-Driven VR Worlds in an Elementary School Curriculum

- The solar system, created at the request of a fourth-grade teacher, allows the students to see both the moving solar system in scale (with the sun and planetary diameters in the same scale as their orbits) to get a feeling for the vast scale of the distances involved and an enhanced view with the planets and their moons enlarged to allow the students to investigate the similarities and differences between the planets. The sun, planets, and major moons rotate and revolve at a rate where one second in the virtual world equals one day in real life.
- The human heart, created at the request of Kathy Madura, the physical education teacher, is an animated beating heart with arrows showing how the blood flows through its various chambers. Optional transparency allows the students to see the 3D nature of the blood's passage through the heart, and a movable clipping plane allows easier viewing of the internal structure.

USAGE

The teachers are responsible for structuring the learning experience, and they guide the actual sessions using the models. Typically, we initially operate the controller for the teachers, but as they gain experience, they take over the controller themselves.

A typical pattern of usage can be seen in a recent experience with the heart model. The physical education teacher had previously used Heartland, a model of the human heart realized as a collection of kid-sized carts, pulleys, balls, hoops, and tunnels laid out on a gymnasium floor. The kids wheeled about on the carts, picking up corpuscles and moving from one chamber to another.

This teacher used the heart QuickWorld to reinforce that earlier learning. The teacher ran sessions of approximately 10 minutes with three fifth-grade students at a time. She began each session by reminding the students of Heartland and relating the anatomical elements of the heart (chambers, valves, aorta, etc.) and the animation of blood flow to that experience. She often began by relating functionality played out in the Heartland model (for example, oxygenation) and seeing if her students could relate that process to the QuickWorlds model. She used about a half dozen different clipping planes during each session. Students readily agreed when directly asked whether the model has improved their understanding; not unexpectedly, the students were impressed with the technology.

EXPERIENCE AND FUTURE WORK

The students are certainly excited about this addition to the curriculum, and we continue to get many requests from interested teachers (for example, adding lungs, cholesterol, leaky valves to the heart model; atoms; more insects.) In future semesters, as we continue the QuickWorlds work, we will be creating more models, and adding more details to the existing models.

It is still too early to give any conclusions about whether these worlds are effective. We believe that the appropriate path this year is to generate a variety of QuickWorlds based on the teacher's requests, make them available to the teachers and their students, and see which

ones the teachers feel are most effective. This experience will then drive the kind of QuickWorlds and their attributes developed in the second year, with the eventual goal of empirically investigating learning effectiveness relative to unsupplemented instruction and non-immersive interfaces. The novelty effect is also very strong, as the children are very excited ("sweet!" was a common exclamation when first seeing the virtual heart) about this new technology. This has advantages and disadvantages. The children are excited by the experience but may not be focused on the lesson to be learned. We would like to reach the point where the ImmersaDesk hardware loses its novelty for the children, and the visualizations themselves become the focus.

We currently have at least one person on hand at the school for one day per week to run the VR equipment for the teachers, and talk to them about any new QuickWorlds requests. This will expand to two days a week in the next school year. We are in the process of making the models from the existing QuickWorlds available in VRML-1 format on our Web pages, so other educators can download them and use them in more conventional VRML browsers.

ACKNOWLEDGEMENTS

This research was made possible through major funding from the National Science Foundation, specifically EIA-9802090, EIA-9720351, and DUE-9979537.

The virtual reality research, collaborations, and outreach programs at the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago are made possible by major funding from the National Science Foundation (NSF), awards EIA-9871058, ANI-9712283, ANI-9730202, and ACI-9418068, as well as NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement ACI-9619019 to the National Computational Science Alliance. EVL also receives major funding from the US Department of Energy (DOE), awards 99ER25388 and 99ER25405, as well as support from the DOE's Accelerated Strategic Computing Initiative (ASCI) Data and Visualization Corridor program.
www.evl.uic.edu/tile

References

1. Johnson, A., Moher, T., Ohlsson, S., Gillingham, M. (1999). The Round Earth Project: Collaborative VR for conceptual learning. In *IEEE Computer Graphics and Applications*, 19 (6), 60-69.
2. Johnson, A., Moher, T., Ohlsson, S. (1999). The Round Earth Project - Collaborative VR for elementary school kids. In *SIGGRAPH 99 Conference Abstracts and Applications*, 90-93.
3. Moher, T., Johnson, A., Ohlsson, S., Gillingham, M. (1999). Bridging strategies for VR-based learning. In *The Proceedings of CHI 99*, 536-543.
4. Moher, T., Johnson, A., Lin, Y., Cho, Y. (2000). Systematic observation in a synthetic environment. In *The Proceedings of the Fourth International Conference of the Learning Sciences*, June 14-17, 2000.
5. Cuban, L. (1986). *Teachers and machines: The classroom use of technology since 1920*. New York: Teachers College Press.
6. Herman, J. (1994). Evaluating the effects of technology in school reform. In B. Means (Ed.), *Technology and education reform: The reality behind the promise* (pp. 133-167). San Francisco: Jossey-Bass, 1994.

A shoestring arts laboratory in Brazil uses a pioneering educational approach and brings social commitment to computer arts. This paper discusses its background and achievements.

The Electronic Arts Unit (Núcleo de Arte Eletrônica, NAE) at Pontifício Universidade Católica do Rio de Janeiro is an arts laboratory where educators, graduates, and undergraduates from science and arts disciplines come together to create and produce art and design projects. At all stages in the projects, a methodology of dynamic, democratic teamwork prompts participants to play new roles: musicians choose colors, animators discuss codes, and programmers create images. Clients range from God to the Devil. We design rock band sites and interactive Bibles! Research projects cover social issues like new media literacy (which is ushering in a new society of netizens, net-fringers, and outsiders), Web design and its cultural implications, and designing ATM interfaces for users with different levels of schooling. Educational results prove we are on the right track: student participants improve in terms of knowledge and skills; develop high levels of creativity, maturity and social commitment; and are now ready to help shape our unequally connected society into something much better for all of us.

NAE: A LOCUS FOR CREATIVE USE OF COMPUTER-BASED TECHNOLOGIES

Until 1992, we were just a bunch of enthusiastic computer graphics educators with different backgrounds (in art, design, computer science, video, and photography) willing to work as a group on innovative multimedia educational projects. We discussed teaching methods, strategies, and course content, but mostly we shared complaints about the enormous workload involved in keeping up with recent hardware and software developments.

In spite of our complaints, we still found time to dream. And we dreamed of having a lab for creative research and development; a place where we could experiment freely on new ideas and techniques — an educational space where we would prepare our students to perceive the real challenges of our digital society and fight to improve it with their spirit of curiosity, their knowledge, and, above all, their social commitment.

The Electronic Arts Unit was set up in mid-1992. All we had was a graphics workstation donated by Sun Microsystems and some old furniture, but that seemed to be more than enough. After all, in those early days, what other educator in Brazil had research and development facilities, an Internet connection, an art domain, plus the valuable help of four student grantees and a group of incredible people to work with?

Our idealistic, pioneering spirit, plus our disparate backgrounds, led us to adopt a methodology of dynamic, democratic teamwork. Although each of us had a basic pre-defined role, according to our core expertise, we encouraged all project participants to play other roles at different stages of the project. At our meetings, musicians would have a say in the choice of colors, animators would end up discussing codes, and programmers had to help create images, whenever possible, necessary, or desirable.

We used to call ourselves The Brancaleone Army (after the 1960s Italian comedy film about a group of medieval bunglers on their way to the Crusades). Individually frail, we were made strong by acting courageously as a group. Like even the smallest pieces of a jigsaw puzzle, each of us was equally important and indispensable in bringing the project to completion. There was a feeling of optimism, and a sense that, regardless of all our differences, we were in this together, and we would open up new and enthralling horizons for our students.

INNOVATIVE - OUTRAGEOUS? - IDEAS

Working on this philosophy and methodology, we then started to look around for projects that might lend themselves to innovative experimentation. And we soon found them.

The Brazilian Research and Development Council (CNPq) gave us a grant to investigate the use of computing as a teaching aid for various subjects in arts and design curricula, and as a tool for facilitating relations among the disciplines offered in these areas. Experiments have involved educators and groups of students in subjects like free-hand drawing, ceramics, mechanical systems, and visual language. By setting up a hybrid laboratory, where various techniques, equipment, and materials could be handled in parallel, we provided these educators with a combination of concrete, intuitive techniques and abstract, logical methods at various stages of the process of artistic creation. For instance, computer graphics and ceramics classes were taught (on an experimental basis) at the same hybrid lab. With the help and orientation of our research team, students and teachers were encouraged to exchange, and share, their tools, theories, and techniques to accomplish common tasks. As a result of working on both clay and computers at the same time and in the same place, outrageously original, creative ideas surfaced. In one of the experiments, students molded a shape in clay, then stuck the wet clay shape onto the monitor screen and used it as a template for plotting and modeling the virtual object (see Figure 2). We called this project “Dirty hands on the keyboard!”¹

Word soon got around about our working philosophy and results, and interesting projects came knocking on our door. When the top Brazilian Rock band Barão Vermelho came into the office wanting an interactive track for their new album (for release three months later!), we did not have a clue how much to charge, much less how to do it, but we accepted the challenge. They gave us complete freedom to create whatever we wanted (and later confessed that they had no idea what you could do with an interactive track either!). All they had was a copy of the Rolling Stones’ “Stripped” album, and they just asked us to do a better job, as if our rag-tag army could outdo the Stones’ big-time productions!

Over sleepless nights spent worrying about how the whole thing would turn out, the basic concept of the project began to emerge. What do people do when they get home after buying a new album? They kick off their shoes, flop onto the sofa, and go with the pleasurable flow of the senses. So what we had to do to make the project a success was to create auditory, visual, and cognitive sensory pleasure. The navigation flow of the interactive track had to be as intuitive, as pleasant, and as free-form as the audio flow was. No imperative commands (“go to the next page!”), no standard icons and buttons, but suggestions, surprises, tricks, and fun.

Teaching Computer Graphics in Brazil: Social Commitment, Creativity and Passion—Against All Odds!

Its non-directive interface, its random situations, and the many surprises in store for users on the way helped give the project its intuitive atmosphere, its nonlinear, non-repetitive navigation structure. It is quite impossible to go through it the same way twice! The biggest surprise is when Frejat, the band leader, talks to the user. With his familiar hoarse voice and casual style, he will sometimes give tips and clues, point to hidden tricks or just tease users into clicking on danger spots, all for the purpose of destroying their preconceived notions of standard navigation. We asked Frejat how he would say something like: “Go to the next screen.” He laughed, and said: “Only crabs go backwards.” We recorded a lot of these unusual, good humored commands and used them throughout the project.

We created a screen for each of the audio tracks, designing them always as a complement to the music. The screen either adds a new dimension to the song or creates a dynamic conflict. For the track “Só as maes são felizes” (Only mothers are happy), we used the image of a dirty toilet bowl, as depressing and lowlife as the lyrics (in spite of the song’s title). You have to flush it to watch the video (see Figure 1). Unfortunately, the CD only allowed us about 100Mb for this track. It was like having a short blanket: whenever we tried to cover our noses, our feet would be showing. Three months later, it was ready. The first rock CD with an interactive track developed and produced in Brazil, and it sold more than 150,000 copies in the first couple of weeks! We never expected that sort of success, and, of course, had not charged accordingly. But we do have the gold and platinum record trophies (for more than 400,000 copies sold) hanging on our wall!

This was at the onset of the multimedia “boom,” and, in terms of the Internet, we were still just crawling. The magic of the Web was just starting to enchant and entice us. With a little bit of imagination (perhaps a lot), HTML code could be seen as a new graphic toolset for artists and designers. The Web’s “chameleon-like” potential for constant change and updating, and its worldwide reach, were definitely new challenges for us. In those early days, we were free to create entirely new things. There were no patterns to follow or models to copy in the pre-dawn of Web design. On the other hand, technical constraints were far more limiting to artists. We had to be really imaginative to perceive the artistic potential of the crude tools and techniques then available and make creative, intelligent use of them. At NAE, we have been working on pioneering, innovative Web design projects since 1994, and by publicizing our work and circulating our ideas, we have contributed to training a major contingent of professionals in this field in Brazil.

Being a shoestring, academic arts lab, and having university support to cover our basic running costs, we were always able to afford the luxury of choosing which projects we wanted to pursue. As a result, we work with a wide range of partners and clients. We have devised Web-based and multimedia projects for both commercial institutions and non-profit organizations, in Brazil and other countries (Figures 1, 2). Our motto is “never do what you already know how to.” Each project must raise a new challenge, something to be researched, learned, taught to students and professors, and spread out into the community.

Over the years, our work has gained increasing visibility. We are often invited to present our ideas and projects on TV programs dealing with art and technology issues. Our opinions and achievements very often get space in the major national press, and we are often asked to participate in debates, juries and meetings relating to the arts and design.

Many of the educators who were part of the team in the early years still work with us. Many students who have already left, and are now professionals working at other institutions, still take part in many of our project meetings, either as guests or as active collaborators. We still share our original feeling of optimism. Looking back, we realize we really did open up new horizons for our students, but we still have a lot more to do.

SOCIAL COMMITMENT AND RESPONSIBILITY: ELECTRONIC ART FOR THE REAL WORLD

Two months ago, I came home to find my 10-year old daughter anxious to tell me what she learned at school that day. As she combed my hair with her small, delicate hands, she gave me a dramatic report on child labor in our country, where almost four million children from the ages of five to 14 are exposed to physical, psychological, and social hazards, and to intolerable forms of exploitation, often working long hours in unsafe, unhealthy conditions for very little pay. Her teacher had told them about the many children who have to carry 50-pound sacks of oranges on their backs to earn less than 50 cents a day. And about the thousands of children who work harvesting sugar cane, who often suffer serious cuts from the enormous knives they have to use to cut the cane with their small, delicate hands. She then looked deep into my eyes, her own eyes filled with tears, and asked me how she could possibly help change that situation. I told her: “You care. That’s already the first and maybe the most important step.”

Some years ago, I questioned the role of an electronic artist in a developing country. What is the point of producing art using expensive, sophisticated technologies in a context of extreme poverty and social inequality?²

In parallel with the projects for commercial or institutional use that we pursue at NAE, we also run a series of research projects on social issues raised by the emerging Internet society. Sponsored by research and development agencies and other public sector organizations, these projects are concerned with the relationship between Web design and economic, social, and cultural diversity, and with new media illiteracy in developing countries. Our project proposal (Internet, Illiteracy and Social Exclusion), which has recently been awarded the very special government grant Cientistas do Nosso Estado (Scientists of our State), focuses on the emerging actors inside, at the fringes of, or outside the Net society. As part of the project, we are currently designing a site to show how people living in different socio-economic circumstances in Rio de Janeiro perceive and understand the Internet, how it affects their lives, and its implications for their future.

No doubt one of the main challenges our digital world faces today is actually to incorporate the communities and groups sidelined by the connected society. “Only a fraction of people in the world have a presence in cyberspace; the rest are outsiders,” says Hertz.¹ The

expanding use of computers in developing countries, alongside dramatic levels of poverty, deprivation, social exclusion, and violence, calls for careful analysis of the social and cultural specifics in each context, and discussion of the computer's impacts on less economically and technologically favored societies.⁵

Today, I consider it crucial to teach computer graphics in Brazil, not just for the sake of technological development and progress (there is no such neutral thing) but to improve the quality of life in our country. Although most of those who have access to high-end resources, including my students, belong to the upper layers of the economic pyramid, computer graphics is perhaps the main channel through which Brazilian artists and designers are discussing human needs and fighting for human rights.

AGAINST ALL ODDS

Seen from where we stand today, things looked a lot easier for us (computer graphics educators) before the astounding, overwhelming Internet boom. There were not as many plug-ins to download, techniques and innovations to learn, and digital newsletters and manuals to read as today. The information at our disposal was not so massive, relentless, and alluring. But, above all, there were fewer discussion lists and messages to reply to, every single day, and less guilt about when you just can't! Our days are much longer, if you include all those extra working hours we need to stay logged on at home to keep abreast of all the enticing information. And increasingly, our students know a lot more techniques than we do, for sure!

But the essence, the very nature of our educational practice has changed very little, or perhaps not at all. Overloaded and burned out as we always are with students' questions, assignments, meetings, and the constant need to stay current,² it can be quite easy to forget

that basic, most fundamental reason for teaching: to prepare people to develop new ways to help bring greater quality of life to the world. That is what we, as educators, really need to focus on.

The Electronic Arts Unit is still located at that same small office (and we will remain there as long as our computers fit in the room!). We still see ourselves as a Brancalone Army. But the unit is well known today as a special place where you will find interesting people, where you can experiment on new techniques and get information on contests, exhibitions, and conferences. It's a place where your opinions are welcome, and, most importantly, where you can let your ideas flow (because there's always someone there who believes in them).

Against all odds, it is a place where people learn how to cultivate their spirit of curiosity, improve their knowledge, and develop strong social commitment. Just the way we dreamt...

Notes

1. The results were presented at several conferences in Brazil, at SIGGRAPH 95, at ISEA'95 (Canada), and at SIGRADI'99 (Uruguay).

References

1. Hertz, P. (1997). Colonial ventures in cyberspace. *The Leonardo Gallery, LEONARDO*, 30 (4), 249.
2. Rubin, C. B., Weintraub, A. & Pointdexter, D. (1994). Faculty guidelines in computer-based media in fine art and design. *Computer Graphics*, 28 (4), 264-267.
3. Spitz, R. (1993). Qualitative, dialectical and experiential domains of electronic art. *The Art Factor. Fourth International Symposium on Electronic Arts, November 1993*, 161-166.
4. Spitz, R. (1995). Computers in arts and design education: Impregnating the digital world with texture, dust, and noise. *Computer Graphics, August 1995*, 16-18.
5. Spitz, R. (1999). Internet, the WWW & human communication: a new Tower of Babel? In Couto, R.M.S. & Oliveira, A.J. (Eds.), *Formas do Design: por uma metodologia interdisciplinar* (pp. 103-127). Rio de Janeiro:2AB Editora Ltda.

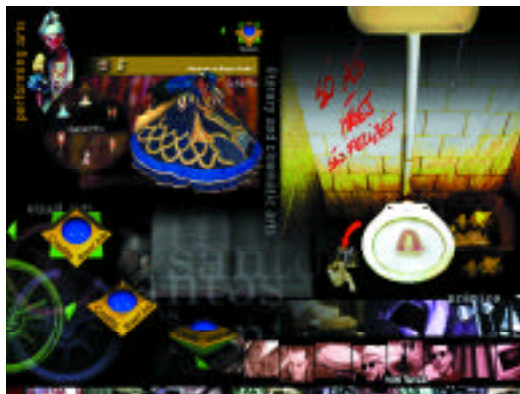


Figure 1. Examples of NAE projects (for Warner Records, Bar o Vermelho, University of Florida, Brazilian Foreign Ministry, Edições Loyola, and INFNET).



Figure 2. Examples of NAE projects on social and academic issues.

Over the past few years, computer graphics courses in computer science and engineering programs have become more programming oriented, primarily due to the availability of the OpenGL API on virtually all platforms. OpenGL is easy to learn and close enough to the hardware that it can support programming applications and understanding of graphics concepts and the underlying algorithms.

Although this change has led to far more exciting classes than the old survey or algorithms-based classes, for many instructors there is a significant startup problem. This is especially the case in the many schools where there is no “expert” on computer graphics available to answer questions and help instructors with their first programs. This tutorial should help those interested in using OpenGL for their courses to learn more about its capabilities and become familiar with two possible approaches to an OpenGL-based course. Hopefully, the participants will also form a community that can support each other and other new users of these approaches.

The presenters are authors of two of the most widely used computer graphics textbooks for computer science and engineering students. Both have made the switch from other approaches to ones using OpenGL.

This tutorial provides a short introduction to the main features of the OpenGL API and then focuses on two very different ways it can be used to teach a senior-level computer graphics course. Participants will be able to try some sample programs from the instructors' courses. All programs will be available to the participants.

Participants will see two very different types of classes that go beyond the difference in programming languages. Angel's approach stresses the integration of computer and traditional visualization, and leads to many applications in scientific visualization. Hill's approach devotes more time to modeling and visual realism, and incorporates more material on fractals.

SPEAKERS

Edward Angel is a professor of computer science and electrical and computer engineering at the University of New Mexico. He has been teaching and doing research in computer graphics, image processing, and scientific visualization for over 25 years. At UNM, he is the first Presidential Teaching Fellow. He is the author of “Interactive Computer Graphics, A Top-down Approach with OpenGL” Second Edition (Addison-Wesley, 2000), which has been adopted by over 200 schools. He has taught over 100 professional short courses worldwide, including at the annual SIGGRAPH conference.

Sandy Hill is a professor of electrical and computer engineering at the University of Massachusetts, Amherst. He has been teaching and doing research in computer graphics, digital signal processing, and geometric modeling for 29 years. He is the author of “Computer Graphics, with C++ and OpenGL” Second Edition (Prentice Hall, 2000). He has taught short courses extensively through the National Technological University (NTU) and in industry. He has received the Outstanding Teaching Award from NTU for the last 10 consecutive years.

INTRODUCTION

“In order ... to appreciate the requirements of the science [of electromagnetism], the student must make himself familiar with a considerable body of most intricate mathematics, the mere retention of which in the memory materially interferes with further progress.”

JAMES CLERK MAXWELL³

Classical electromagnetism is a fundamental underpinning of a technical education, but one of the most difficult subjects for students to master. It is also a subject in which mathematical complexity quickly overwhelms physical intuition. With the aid of animation, we are developing a treatment of electromagnetism designed to help students develop intuition about the dynamics of electromagnetic phenomena, in a manner independent of advanced mathematics.

Laurillard² cites two key criteria for selecting subjects in the college curriculum for this kind of treatment. To justify the extensive resources required for technological development, the subject must be: widely taught and widely acknowledged to present difficulties for students. In the standard science and engineering curriculum, classical electromagnetism satisfies both of these criteria. It is widely taught and also widely misunderstood. Why is this? It is because students have few pre-existing models for electromagnetic phenomena or of the concept of fields. Since much of our learning is done by analogy,⁵ students have a hard time constructing conceptual models of the material they are trying to absorb. The standard textual approach to teaching this subject does little to help students establish such conceptual models, because a purely textual approach does little to connect the dynamics of electromagnetic fields to the student's everyday experience.

However, there is a way to make that connection for many situations in electromagnetism — an approach that has been known since the time of Michael Faraday, the inventor of field theory. Faraday was also the first to understand that the shape of field lines is a remarkable guide to their dynamics — a guide that does not require use of advanced mathematics to understand. By trial and error, Faraday deduced that field lines exert a pull parallel and a push perpendicular to themselves. Knowing the shape of field lines from his experiments, he was able to understand the dynamical effects of the fields based on simple analogies to ropes and strings, without recourse to advanced mathematics. It is this Faraday approach that we are pursuing, with the goal of helping students gain intuition about electromagnetic dynamics. Our means to achieve that goal is animation.

The following question arises immediately: If animation is an effective way to develop student intuition, why is it almost never used in introductory courses in electromagnetism? The answer is that this method is of only modest use if the field lines are displayed as static images. However, the power of the method increases dramatically if the field lines are animated. The mind has an enormous capacity to integrate

time-changing visual information into a coherent dynamical whole — a capability that evolved because it is fundamental to survival. With animation, one can appreciate the effects of the stresses transmitted by fields in an immediate and visceral way, by watching how things evolve in time in response to these stresses.

Animation has not been used to display electromagnetic fields in the past (with notable exceptions) due to the twin difficulties of producing such animations and then delivering them to the student in an easily accessible manner. The enormous increase in computing power over the last decade, and the advent of the World Wide Web, has made both production and delivery of animations, integrated into textual development, an increasingly viable proposition. We are taking advantage of these technologies to implement Faraday's insights, using video clips to display actual experiments, as well as producing computer visualizations and animations of the electromagnetic field lines in those experiments. These visualizations of the field lines make the unseen seen, so that students can come to an understanding of what is happening dynamically, via analogies to familiar concepts.

EXAMPLES OF VISUALIZATIONS

“To understand this point, we have to consider that a [compass] needle vibrates by gathering upon itself, because of its magnetic condition and polarity, a certain amount of the lines of force, which would otherwise traverse the space about it.”

MICHAEL FARADAY¹

“It therefore appears that the stress in the axis of a line of magnetic force is a tension, like that of a rope.”

JAMES CLERK MAXWELL⁴

Contrast the usual way of explaining the torque on a compass needle in a background magnetic field with Faraday's approach to understanding the same phenomenon. In the standard explanation, we appeal to the notion of atomic currents in the needle, circulating in a plane transverse to its dipole axis. We then consider the torque on such a current loop. We usually take a rectangular loop of wire carrying current i in a background field \mathbf{B} . We look at the various forces on the sides of the rectangular loop to deduce the net torque on the rectangular loop, which tends to align the compass along the background field. The advantage of this procedure is that it yields a quantitative calculation of the torque. The disadvantage is that the explanation requires several relatively abstract steps, which most students cannot reproduce in any coherent fashion. Thus, although they memorize the result, students subsequently have little intuitive feel for why it should be so.

In contrast, consider how Faraday explained the torque on a compass needle, and thus its oscillations. First, he used his intuition about the shape of field lines based on his experiments with magnets and iron filings. He then appealed to the concept of a pull along the field line

to infer the dynamical effects associated with that field configuration. In the case of the compass needle in a background constant field, he drew a field configuration for the sum of the magnetic field of a dipole whose dipole vector makes an angle to the vertical, plus that of a constant vertical field. Faraday then understood the oscillation as due to the tension in the field lines pulling the needle into alignment with the background field, with the needle then overshooting.

An animation of this behavior makes the oscillation seem natural and intuitive. We argue that both of the above explanations should be provided to the student. The first is quantitative and appeals to students who are analytical in their thinking. The second is qualitative (although it can be made quantitative) and much more intuitive, and it is comprehensible to students of all persuasions, because it can be understood by analogy to concepts they already have. Our contention is that one year after taking a course in electromagnetism, average students will not remember the details of the first explanation. However, if they have “seen” the second, they will continue to have a mental model as to why compasses “work” this way.

THE FALLING MAGNET EXPERIMENT

One of the main thrusts of our effort is to make the unseen seen, to use the power of sophisticated 3D animation to show the student “actual” phenomena, with the ability to add to that visualization things that we cannot ordinarily see. As an example of this, consider an experiment, and then a virtual recreation of that experiment.

A magnet falls in a Plexiglas tube through a conducting copper ring. When the magnet approaches the ring from above, eddy currents are set up so as to prevent an increase in the flux through the ring, and the magnet undergoes a clear deceleration. Figure 1 is a frame from a 3D animation of this motion. As the magnet moves toward the ring from




Figure 1: The magnetic field lines of a magnet falling under gravity through a conducting non-magnetic ring (for example, a ring made of copper). The field lines shown are those of the total magnetic field (the dipole field of the magnet plus the magnetic field due to the eddy currents generated in the ring, which are, in a sense, such as to try to keep the flux through the ring constant).

above, it is repelled by currents induced in the ring. In the animation, the field is compressed as the magnet falls toward the ring, and the slowing down is naturally interpreted by the mind’s eye as a deceleration due to an upward push due to the compression of the field. When the magnet falls through the ring and is below it, the eddy currents reverse direction so as to now prevent a decrease in the flux through the ring, and the magnet is again decelerated. In the animation, the field is now stretched out as the magnet falls away from the ring, and the slowing down is naturally interpreted by the mind’s eye as a deceleration due to the upward pull of the stretched field lines. This overall animation is a good example of our approach. The visual treatment does not replace the traditional explanation, but complements and expands on it, and in a way that makes “intuitive” sense to the student.

MAGNET BEING PULLED AWAY FROM A COIL OF WIRE.

Let us consider a final example of our approach. Consider an experiment in which a magnet is moved along the axis of a coil of wire. Initially, the magnet is at rest close to the coil, and then is pulled away from the coil at constant speed and brought to rest farther away from the coil. As the magnet is pulled away from the coil, an ammeter registers current in the coil in a direction such as to try to prevent the decrease of flux through the coil. The sense of that current is such that the coil and the magnet are attracted (the agent moving the magnet must do work to pull the magnet away from the coil).

Now consider a computer visualization of this process. Figure 2 is one frame of an animation of this process, at an instant of time just before the magnet comes to rest. The field lines have a hard time “getting through” the coil, since the sense of the current in the coil is such as to try to keep the number of field lines threading the coil from decreasing. Thus the field lines get “hung up” on the coil as they try to move through it. The intuitive sense that one gets in




Figure 2: The magnetic field lines of a magnet being pulled away from a coil of copper wire. The field lines shown are those of the total magnetic field, (the dipole field of the magnet plus the magnetic field due to the eddy currents generated in the coil). We show many field lines in latitude, instead of only one field line in latitude, as in Figure 1. The field is symmetric about the axis of the coil.

watching this animation is that the agent moving the magnet must do work to pull the field lines “through” the coil. We emphasize that these animations are based on quantitative calculations. These are not cartoons.

MOVING FIELD LINES

The concept of moving field lines is unfamiliar to many professional scientists and engineers. Since this concept is central to our approach, we discuss it here, for the benefit of that audience.

The magnetic field lines above are defined in the usual way: a field line is everywhere tangent to the local field. We make no attempt to have the density of field lines correspond to field strength (this is impossible in 2D projections of 3D fields in any case.⁷ How do we define the motion of magnetic field lines in the above animations? Consider the following thought experiment.

We have a solenoid carrying current provided by the emf of a battery. The axis of the solenoid is vertical. We place the entire apparatus on a cart and move the cart horizontally at a constant velocity \mathbf{V} as seen in the laboratory ($V \ll c$). We intuit that the magnetic field lines associated with the currents in the solenoid should move with their source. How do we make this intuition quantitative? First, we realize that in the laboratory frame there will be a “motional” electric field $\mathbf{E} = -\mathbf{V} \times \mathbf{B}$. We then imagine placing a low-energy test electric charge in the magnetic field of the solenoid. The charge will gyrate about the field, and the center of gyration will move in the laboratory frame because $\mathbf{E} \times \mathbf{B}$ drifts in the $-\mathbf{V} \times \mathbf{B}$ electric field. This drift velocity is just \mathbf{V} . That is, the test electric charge “hugs” the “moving” field line, moving at the velocity our intuition expects.

In the more general case (for example, two sources of field moving at different velocities), the motion we show in our magnetostatic computer visualizations has the same physical basis. It is the drift motion we would observe for hypothetical low-energy test electric charges initially spread along the various magnetic field lines, drifting in the electric field that arises because of the time-changing magnetic field via Faraday’s Law.

SUMMARY

“...a simple precis is that these improvements are attempting to nurture a sense of wonder among students about the natural world, and to maintain students’ active curiosity about this world while equipping them with tools to explore it and to learn.”

SHAPING THE FUTURE: NEW
EXPECTATIONS FOR UNDERGRADUATE
EDUCATION IN SMET⁶

Let us return to the objectives of our approach. One of our primary aims, an aim that is fulfilled if we are careful in what we present, is to engender a sense of wonder in the student. The 3D visualizations that we have created and plan to create are visually compelling. They engage the student’s imagination because they show the world in

a photo-realistic way, including representations of phenomena that heretofore could only be seen in the mind’s eye. In large lecture courses in the freshman year, one of the purposes is to inspire students to invest the time to pursue quantitative mastery of the subject outside of lecture. Our extension of the pedagogy in this subject will be successful in large degree as long as it arouses interest and excitement by engendering a sense of wonder.

Beyond engendering a basic sense of wonder, what is the central student learning need that we are trying to meet? It is this: Students need an enormous amount of help in understanding the nature of fields. The central learning objective of introductory courses in electro-magnetism is to help students understand how fields are generated, how they mediate the interaction of material objects, and how they propagate. Our contention is that in the standard pedagogy, this learning objective is not well fulfilled. Our approach to help remedy this deficiency is to give the fields a more prominent role in the pedagogy, by literally making them more visible. They are thereby made more understandable dynamically, based on students’ pre-existing models of the behavior of strings and rubber bands.

The use of animated visual displays of field lines has many advantages. They continually remind the student that it is the field that mediates interactions between material objects — that the field has as much “reality” as the objects themselves. By stressing the pushes and pulls transmitted by the fields, we stress the importance of the fields themselves as the mediator of interactions. Ultimately, animations allow students to understand intuitively what is happening dynamically simply by looking at the shape of the field lines, once the eye and the mind are trained to this purpose. It is this intuition that we seek to develop.

ACKNOWLEDGEMENTS

This work is supported by National Science Foundation Grant 9950380, an MITClass of 1960 Fellowship, The Helena Foundation, the MIT Classes of 51 and 55 Funds for Educational Excellence, the MITSchool of Science Educational Initiative Awards, and MITAcademic Computing.

References

1. Faraday, M. (1855). *Experimental researches in electricity, Volume III* (p. 227). London: Taylor and Francis (Reprinted by Dover Publications, New York, 1965).
2. Laurillard, D. (1993). *Rethinking university teaching: A framework for the effective use of educational technology* (pp. 181-182). London: Routledge.
3. Maxwell, J. C. (1855). On Faraday’s lines of force. *Transactions of the Cambridge Philosophical Society, X* (Part I). (As quoted in Simpson, T.K., *Maxwell on the Electromagnetic Field* (p. 55). New Brunswick: Rutgers University Press, 1997).
4. Maxwell, J. C. (1861). On physical lines of force. *The Philosophical Magazine, XXI*. (As quoted in Simpson, T.K., *Maxwell on the Electromagnetic Field* (p. 146). New Brunswick: Rutgers University Press, 1997).
5. Redish, E. F. (1996). *New models of physics instruction based on physics education research*. Invited talk presented at the 60th meeting of the Deutschen Physikalischen Gesellschaft, Jena, March 14, 1996.
6. *Shaping the future: New expectations for undergraduate education in SMET*, a report on its review of undergraduate education by the Advisory Committee to the NSF Directorate for Education and Human Resources, 1996.
7. Wolf, A., Van Hook, S.J., Weeks E.R. (1996). Electric field line diagrams don’t work. *American Journal of Physics, 64*, 714.

Contents

75	4 + 4 Fingers Direct Manipulation with Force Feedback	88	Medieval Chamber
76	ActiveCube	89	MicroTelepresence
77	Augmented Groove: Collaborative Jamming in Augmented Reality	90	Musical Trinkets: New Pieces to Play
78	Autostereoscopic 3D Workbench	91	Muu: Artificial Creatures as an Embodied Interface
79	Autostereoscopic Display for an Unconstrained Observer	92	Networked Theater: A Movie Production System Based on a Networked Environment
80	CYPHER: Cyber Photographer in Wonder Space	93	Panoscope 360°
81	Danger Hamster 2000	94	Plasm: In the Breeze
82	Gait Master	95	Retinal Direct Imaging
83	HoloSpace	96	RV-Border Guards: A Multi-Player Entertainment in Mixed-Reality Space
84	InTheMix	97	VaRionettes
85	Jamodrum Interactive Music System	98	V-TOYS: Visually Interactive Toys
86	LaserWho	99	X'tal Head: Face-to-Face Communication by Robot
87	Magic Book: Exploring Transitions in Collaborative AR Interfaces	100	You Were There

EMERGING TECHNOLOGIES

J. JEFFREY CLOSE



Committee

CHAIR

J. Jeffrey Close

SUBCOMMITTEE

Christopher Carney
Callison Architecture Inc.

Sharon Johnson
Pacific Northwest National Laboratory

Cici Koenig
California Institute of Technology

Richard May
University of Washington

Preston Smith

JURY

David Ebert
University of Maryland

Steven Feiner
Columbia University

John Funge
Sony Computer Entertainment America

Andrew Glassner
Consultant

Mk Haley
Walt Disney Imagineering

Jan Hardenbergh
Real Time Visualization/Mitsubishi

Mary Lou Jepson
Periscopic, Inc.



WELCOME TO SIGGRAPH 2000 EMERGING TECHNOLOGIES: POINT OF DEPARTURE!

This program presents new technologies in several areas of computing research. Of course, Emerging Technologies has always been about fun, hands-on interactive projects, where you can walk through, wave your arms, make faces, talk to computers, and push buttons. These installations are definitely fun, but they also generate reflection. It is the year 2000, a subject of fiction, jokes, and prophecies. If the past century is any guide, and trends continue, astounding changes will occur during the next 100 years. We approach this with a necessary optimism, tinged with responsible skepticism. In SIGGRAPH 2000 Emerging Technologies, we ask you to consider the future in several areas of computing:

- *The Senses.* Input. Output. Display devices, visual rendering, audio rendering, haptics, speech recognition, image recognition, facial recognition, emotion recognition. The technologies for how computers communicate with us, and how we communicate back, define our basic relationship with computing technology. These will define whether computers conform to us, or we to them.
- *Mobility.* Portable computing, handheld computing, ubiquitous computing, smart objects, smart rooms, tangible interfaces. The goal of these research areas is to make it possible for us to use computers wherever we go, even to hide their use within ordinary human activities. Instead of working in the field, the factory, or the office, for the first time we may have an opportunity to carry our workplace with us, or even to eliminate “the workplace” as a special environment. However, we must also insure that this power does not eliminate human freedom, privacy, and intimacy.



-
- *Collaboration.* Multiuser, multiplayer, networked, distributed humans. How do you define your real community? Do you believe there is such thing as a virtual community? The Web, in its present embryonic form, shows signs of promise for connecting people. However, there is also evidence that it can create virtual community while eroding real community. Clearly, if computers are to improve human existence, a central requirement is that they enhance our ability to congregate and communicate and not intensify our isolation.
 - *Robotics.* Is the goal to develop an intelligent, autonomous machine, or a machine that is just intelligent enough to help us? Computers in the form of robots have great potential for performing hazardous work and aiding the disabled. Yet a frightening and recurring theme of science fiction is that we humans may be merely a step in the evolution of intelligent machines. Some scoff at this idea, and some see truth in it and are strident in their reaction.

The choices that we make - the choices that you make - are what will determine that future. We in the technology professions need a Hippocratic oath. As individuals, we must apply our personal ethics to our work. Think about where you want technology to take us, and believe that you have some power in determining that path. If we succeed at this, then one day we may say, as Cicero said with books: "A room without computers is like a body without a soul."


The talented teams at Capstone, Smith Bucklin, AVW, Freeman, and Q LTD, who were indispensable and a pleasure to work with.

4 + 4 Fingers Direct Manipulation with Force Feedback


Makoto Sato
Tokyo Institute of Technology
msato@pi.titech.ac.jp

Unlike similar human interface devices for a direct manipulation environment with force feedback, 4+4 Fingers requires no gloves or heavy, complex metallic structures attached to the user's hands. Users wear nothing more than small, lightweight fingertip devices on four fingers of each hand. They have a realistic sense that their hands are free, and they can make direct use of their hands in the manipulation environment. Without thinking about operating an input device, users


can perform dexterous manipulations with one hand or collaborative movements with both hands as they naturally interact with virtual objects. In this 3D virtual world, users perceive force feedback generated by the device and discover the shape of the virtual object. They can move, rotate, or assemble two different virtual objects at the same time.




Block diagram of the system.



User manipulates virtual objects.



Peg-In-Hole.



Constructed system: SPIDAR-8.

Collaborators
Somsak Walairacht
Keita Yamada
Shoichi Hasegawa
Masahiro Ishii
Tokyo Institute of Technology

With this bi-directional interface, users manipulate physical cubes to construct and interact with 3D virtual objects. All the faces of the cubes are the same and can be connected to any face of any other cube. The computer recognizes the constructed 3D structure in real time, so consistency is always maintained between reality and its corresponding representation in a virtual environment.

Each cube is equipped with a processor for autonomous simulation and communication functions between cubes. The cubes also has a sensor to detect the operational intention of the user and/or a display-actuator output channel. As users interact with ActiveCube in various ways, the displays-actuators present responses. This combination of input and output devices makes the interface intuitive and helps to clarify the causal relationship between the input of the user's operational intention and the output of simulated results.



Children playing with the ActiveCube.

Collaborators
Yuichi Itoh
Michihiro Kawai
Masatsugu Asai
Toshihiro Masaki
Fumio Kishino
Osaka University

Hideo Kikuchi
System Watt Co., Ltd.

Augmented Groove: Collaborative Jamming in Augmented Reality

Ivan Poupyrev
ATR Media Integration and
Communications Research Laboratories
poup@mic.atr.co.jp

Augmented Groove is a musical interface that explores use of augmented reality, 3D interfaces, and physical, tangible interaction for conducting multimedia musical performance. Users can play music together, with or without traditional music instruments, simply by picking and manipulating physical cards on a table. The physical movements of the cards are mapped to changes in musical elements such as timbre, pitch, rhythm, distortion, and reverb. At the same time, users wearing lightweight head-mounted displays can see 3D virtual images attached to the cards. The shapes, color, and dynamics of the virtual images reflect aspects of the music.

In this augmented reality environment, users can see the physical world, virtual objects, and each other, so several users can gather around the mixing table and “jam” together, passing sequences to each other in the same manner in which we pass everyday objects from person to person. In a sense, the music becomes a physical, tangible object, something that can be touched and seen as part of the physical environment.

Musical control in Augmented Groove is deliberately imprecise: performers manipulate short musical sequences, or phrases, rather than individual words. The resulting musical performance is a “computer-supported improvisation” in which the user arranges, mixes, and modifies pre-composed musical phrases on a structural level.



Augmented Groove: augmented-reality musical interface.



Jamming in the Augmented Groove.

Collaborators

Rodney Berry
Jun Kurumisawa
Keiko Nakao
ATR Media Integration and
Communications Research
Laboratories

Mark Billingham
Chris Airola
University of Washington

Hirokazu Kato
Hiroshima City University

*Music Composition
and Production*
Tomoko Yonezawa
Lewis Baldwin

In this reality-enhanced autostereoscopic 3D workbench, viewers without special glasses can perceive 3D images within their reach with little sense of incongruity. A real image is presented in the middle of the

workspace; artificial parallax is narrowed to display 3D objects without interfering with normal motion. Viewers can manipulate virtual objects in the 3D workspace directly, with their own hands.




Figure 1: Sketch of the autostereoscopic 3D workbench system.




Figure 2: Optics of autostereoscopic display. To display virtual object ABC, blue ABC and green ABC are shown in the background, so each ABC can be seen only by the right or left eye, respectively.





Figure 3: Comparison of conventional 3D display and the autostereoscopic 3D display. In our display, a real image of the display is generated in the middle of the workspace and the artificial parallax required to present virtual 3D objects in the workspace is reduced, which leads to less eyestrain and more realistic 3D perception.


Collaborators
 Mitsuo Isogai
 Kenji Suzuki
 Yoshiki Arakawa
 Communications
 Research Laboratory

In this display, observers can freely change position, rotate their heads, and maintain a true stereoscopic view without using special eyewear. Since no fixed parallax barrier geometry could accommodate arbitrary observer position and orientation, the system creates a dynamically varying parallax barrier that continually changes the width and positions of its stripes as the observer moves. Large


stripes would be easily seen by the unaided eye, so the system rapidly animates them in a lateral direction. Each stripe is composed from some number of very slender microstripes, each of which is a switchable LCD element. The result: a stereoscopic display that is continually exact for one moving user, as long as their eye position is tracked.



Phase I




Phase II



Phase III

The three phases of the parallax barrier animation cycle.



One phase shown for a viewer with head turned.

Collaborators

Toto Paxia
Joel Kollin
New York University


Myron Flickner
IBM

CYPHER allows users to build a virtual world based on predefined themes and photograph themselves while they are immersed in it. Users build the world with wooden blocks that represent buildings and objects in the virtual world, and place them wherever they please within a set area. After the virtual world is constructed, users place a human figure block where they want to appear in the virtual world.

Then they instruct a robot to take a picture of themselves in the virtual world. The robot automatically moves or changes its camera angle to find the optimal composition based on its compositional knowledge for the photo of the user in the virtual world. The resulting photo is presented to the user.



Constructing a virtual world

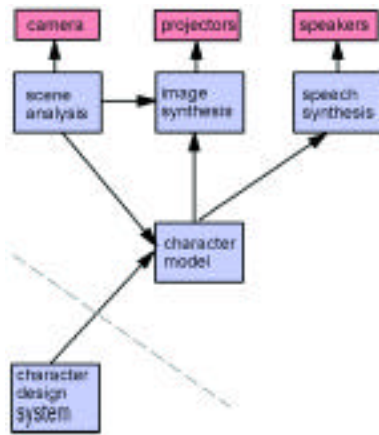


Stages of the Cyber Photographer


Collaborators
Jun Kurumizawa
Keiko Nakao
Jeffrey Chi Yin Ng
Tsutomu Kanegae
Takayoshi Nakato
ATR Media Integration
and Communications
Research Laboratories

Inside a hamsterball, a hamster rolls more or less randomly around an enclosed area. Meanwhile, a camera tracks the position of the ball, and a set of projectors displays a computer-generated face on it. The perceived result is that there is a head rolling around the room, except that the face stays upright. As the ball hits obstacles, avoids collisions, and so on, the face reacts visually and verbally, in accordance with its

character (appearance, speech, moods, attitudes towards objects), which is designed by the audience. Audience members create their own characters, then embody their characters in the hamsterball world. There, the characters respond expressively and in real time to the complex environment.



The DH2K system architecture.




The hamsterball rolls around an enclosed area, bumping into obstacles. It is tracked by an overhead camera, and a character face is projected onto it.



The image projected onto the hamsterball is an expressive face.

Gait Master is a locomotion interface that creates a sense of walking on irregular surfaces. Advanced applications of virtual reality technology often require a credible sense of locomotion. In this project, a two-degree-of-freedom motion platform (one for each foot) is actuated


by two AC motors and chains. The motion platforms track the feet and carry them back to the neutral position. Users can physically walk in a virtual space while their position is maintained, even as they move up and down a virtual staircase.



Overall view of the GaitMaster



Mechanism of the motion platforms




Basic concept of the GaitMaster

Collaborators
Hiroaki Yano
Fumitaka Nakaizumi
University of Tsukuba

Enter the space of digital holograms, the display medium for the new century. The HoloSpace contains research examples of future applications of digital holograms: collaborative design and engineering examples, advertising and commercial applications, and 3D portraits that require no goggles or intrusive devices, and even allow the subject to move!

HoloSpace demonstrates full-color, full-parallax holographic stereograms with a horizontal field of view of 110 degrees and a vertical field of view of 98 degrees. They are as large as 3 meters x 1.2 meters and as deep as 1.2 meters. Viewers look at them from different vantage points: above, below, and inside. Are these 3D displays containing 4D data, or is it the other way around?



Collaborators
Mark Holzbach
Mike Klug
Zach Henson

Imagine standing in the middle of a music mix, quite literally! Your aural environment reacts to your every move, everything you look at, and everything you say. InTheMix employs real-time, digital computation of 3D sound-field models to synthesize an aural environment for presentation via headphones. The aural environment reacts to the listeners' interest, which is inferred from their position and orientation relative to the sounds they hear. The context-dependent content evolves from attention-coaxing sound effects and voices to a selection of production mixes of original multi-track music. Each audio track represents an instrument or vocalist, with an independent dynamic position and orientation in the sound space.

Listeners are allowed to physically roam within a two-meter-radius circle. If they venture past the edge of the virtual space, voices and sounds encourage them to return. While each node of InTheMix may be compelling on its own, several nodes can be linked to form a shared environment for remote participants. Microphones in the headsets allow participants to communicate and share their experiences as if they were physically in the same space.

www.audiosimulation.com



InTheMix allows listeners to explore their imaginations.

Executive Producer
William L. Chapin

Performance
John Tindel
Geoff Rutledge
Lennox Smith
Rachel Wilkinson

Production
Agnieszka Roginska
Lennox Smith
Rachel Wilkinson

Programming
William Chapin
Agnieszka Roginska
Rachel Wilkinson
Hua Zheng
 AuSIM, Incorporated

Parametric Synthesis
Robin Bargar
Hua Zheng
 National Center for
 Supercomputing
 Applications


Inspired by the community spirit of drumming and new technology, the Jamodrum Interactive Music System encourages collaborative experiences with music and real-time computer graphics. Cross-cultural music integration was inspired by the inventors' experiences with communal music making, both in the context of non-Western cultures and contemporary computer music.

The system combines velocity-sensitive input devices and computer graphics imagery into an integrated tabletop surface, where up to three players can participate in a synergetic approach to musical improvisation. Each participant strikes a MIDI pad that controls sounds

and visual imagery to create group compositions. By taking advantage of the players' intuitive abilities and social interaction skills, the Jamodrum embraces both musically trained and untrained participants.

This prototype was the first in a series sponsored by Interval Research Corporation, which ultimately led to installation of a 12-person Jamodrum device at the Experience Music Project in Seattle.

Special thanks to David Liddle, Michael Naimark, Bill Verplank, Becky Fuson, Laurie Anderson, Arthur Hull, and our esteemed colleagues at Interval Research for their support and feedback.



Psychedelic tendencies

*Concept, Software, Hardware
& Interaction Design*

Tina Blaine
Interval Research Corporation

Tim Perkis
Videoscribe

Software & Graphic Design

Greg Jalbert
Imaja

*Audio Engineering
& Sound Design*

Kris Force
Interval Research
Corporation

Hardware Design

Greg Beattie
Tricia Wright
Interval Research
Corporation

Musicians

Carolyn Brandy
Susan Jette
Karma Moffett
Karamba Diabat
Michael Coffey
Bean

Interaction Design

Ignazio Moresco
Hila Dar
Different

Eiko DoEspiritoSanto
Traquitanas

Additional Graphics

Stuart Sharpe
Steve Rubin

Projection

Charles "Bud" Lassiter
Chris Seguine
Interval Research Corporation

In this interactive, gesture-based visualization of the affiliations within a community, users place anchors, each representing a different topic relevant to that community, on the screen. The names of people who are particularly drawn to a topic flow towards the associated anchor. The resulting animation reveals the shared interests and affiliations that comprise the structure of the community.

In the real world, we are continuously (though often peripherally) observing the social patterns around us. In a conference like SIGGRAPH 2000, we notice groups of people who always appear together: students whose common bond we recognize because of their shared sartorial eccentricities, groups of suited businessmen amidst a sea of t-shirt-clad researchers. These observations help us make sense of our complex social environment.

Similar patterns exist in the virtual world, whether as communities of people in an online service, the employees of global corporations, or

electronically registered participants in a conference. Yet in the virtual world, such patterns are usually invisible, hidden within vast databases; they can be seen only if they are explicitly visualized.

LaserWho not only visualizes these social patterns, but does so in a way that recreates (though in a very different way) the fascination of people watching in real life.


www.schwab.com/SchwabNOW/images/Footnote1.gif

www.schwab.com/SchwabNOW/SNLibrary/SNLibog8/Img/SNog8lbml0go.gif

www.neworleans.net:80/entertainment/images/arrow2.gif

www.networksolutions.com/images/home/sm.gif

www.andersen.com/news/images/leftleft.gif



Joe Paradiso
Dana Spiegel
Danah Beard
Ari Adler
Kai-yuh H Siao
MITMedia Lab
Sociable Media Group

Magic Book: Exploring Transitions in Collaborative AR Interfaces

Mark Billingham
Susan Campbell
Human Interface Technology Laboratory
University of Washington
grof@hitl.washington.edu

Magic Book explores transitions between physical reality, augmented reality, and immersive virtual reality in a collaborative setting. A Magic Book looks like a regular storybook with colorful pictures and clear text. When readers look at the pages wearing lightweight head-mounted displays, the pictures pop off the page and come to life as animated objects in virtual scenes. Readers can also fly into the immersive VR world and freely explore the scene. Several readers can gather around

a Magic Book and experience it together. They can view AR scenes from their own perspectives or fly into the immersive world and see each other represented as avatars in the same world. Those who remain in the AR setting have a God's-eye view of their collaborators, who appear as miniature avatars in the virtual scene. Together, readers experience storybook scenes and gain unique perspectives through a variety of navigation styles.



Collaboration within physical space.



Collaboration within AR spaces.



Collaboration within immersive VR spaces.

Ivan Poupyrev
Kazuhiko Takahashi
ATR Media Integration and
Communications
Research Laboratories

Winyu Chinthammit
Duff Hendrickson
Human Interface
Technology Laboratory
University of Washington

Hirokazu Kato
Hiroshima City University

In this demonstration of an alternative interface mechanism for interacting with a graphical environment, real-time video processing tracks multiple stick-like objects in 3D to control on-screen graphical objects. The system tracks three objects, which correspond to a torch,

a sword, and a ball-and-chain (demonstrating point lighting and shadows, reflections, and internal dynamics, respectively). The graphical objects are rendered into an artistic medieval chamber in real time.



The Medieval Chamber real-time rendered environment (torch and ball-and-chain shown).



The user manipulates rendered objects by moving real objects in front of a video camera.

Collaborators
Gabor Nagy
Sony Computer Entertainment

Eric Larsen
Sony Computer Entertainment

Microtelepresence applies telepresence to the microscopic world. A stereoscopic head-mounted display immerses users in a remote microscopic world. Although it is only a few inches square, via microtelepresence, it appears to be about the size of an auditorium. The space is populated with various species of live insects. Relative head movements of the user are sensed with inertial-tracking hardware

and control navigation within the microenvironment. This is accomplished with a three-degree-of-freedom motion stage (basically a converted numerically controlled mill) adapted to a high-quality video microscope. For optimum depth perception, the microscope and display hardware are both stereoscopic.



Navigating a microenvironment.



Collaborators
Dan Gelb
Mike Goss
Mark A. Livingston
John Recker
Hewlett-Packard Laboratories

Lisa Ryner
Stanford University

A swept-frequency resonant tag reader optimized for real-time user interaction and measurement of several continuous parameters enables users to interact with music and graphics in unusual ways through an array of 15 tagged objects. When inserted into the measurement volume (roughly a one-foot cube atop the sensing surface), each object varies a different set of musical parameters or rules.

Since they don't interfere with each another, all tags can be used simultaneously. Every tag measures continuous proximity correlated with angle. Several can be worn around the fingers as rings, creating a virtual velocity-sensitive keyboard. Some objects are embedded with multiple tags so they can also resolve two-axis inclination. Others change characteristics with pressure, enabling a "squeeze" measurement. Some are structured to sit atop the table, launching background sounds and textures, while others can be rolled around, dynamically adding and modifying audio and graphic effects.



Tagged objects, reader, and sensing coil.



Musical trinkets system with generated computer graphics.

Collaborators
Kai-Yuh Hsiao
Ari Benbasat
Massachusetts Institute
of Technology

The primary motivation of everyday conversation is not always to convey messages to others, but rather to reinforce social bonding and coordinate social distance. Muu, an artificial creature with a humorous physical body, was developed to create social bonding with humans. The creature also works as an embodied interface

that mediates the social bonding that people establish in everyday conversations. The mechanism that coordinates the machine-human relationship is implemented using the self-referencing learning classifier system.



Muu: Artificial creatures as an embodied interface.


Collaborators
Shoji Sakamoto
Noriko Suzuki
ATR Media Integration
and Communications
Research Laboratories

Networked Theater: A Movie Production System Based on a Networked Environment

Kazuhiko Takahashi
ATR Media Integration and Communications
Research Laboratories
kylyn@mic.atr.co.jp

In this proposal for a new movie production system that utilizes a virtual environment, a computer network, and the Internet, anyone can be a performer, a producer (director), or an audience member. A motion capture system is connected to a client computer to estimate the performer's movements in real time, and the estimated data

controls a CG character in the virtual environment. At another client computer, a producer creates a movie based on the motion information of the character and the virtual environment, and the movie is distributed through the network to unspecified audiences.



System configuration



Motion capture system



Voice Recognition





Story control panel

Collaborators
Jun Kurumisawa
Shigeo Morishima
Tatsuo Yotsukura
Shinjiro Kawato
ATR Media Integration and
Communications Research
Laboratories

Panoscope 360° is a panoramic viewer that displays a full 360-degree cylindrical image from a single video or data channel. A donut-shaped rear projection screen is placed horizontally above a conic reflector. From within the installation, visitors see a virtual cylindrical image obtained by reflexion.

The installation is designed to simplify both production and presentation of panoramic video programs. It is particularly appropriate for interactive content, where visitors are invited to navigate a space and manipulate elements within it. At SIGGRAPH 2000, Panoscope 360° features Space by Number, the first public presentation of an interactive video program in which visitors are invited to explore a territory by calling numbers displayed in the image.




Panoscope 360° (patent pending)

Concept and design
Luc Courchesne


Illustration
Etienne Desautels

Like the slow, pendulous swings out over the creeks of our youth, Plasm: In the Breeze is very easy to engage; you just hop on and swing. This full body involvement offers a more visceral interaction with the display than the usual “hands-at-a-distance” interface. Visitors can jump on easily and move along after they’ve had enough. When no one is swinging, the “creek bed” cycles through recent and favorite sequences from past performances.

Because the two swings are hung from a connected structure, they interact in ways similar to swings hung on nearby tree branches. This makes each viewer’s experience dependent on the other’s swinging and hopping behavior, which adds to the sense of mutual engagement.



Coupled swings control dynamically generated imagery.




A spiral shows the swinger’s trail fading over time.

Collaborators
Rob Myers
Becky Fuson
Plasmatic Arts


Delle Maxwell

This new stereoscopic display uses a holographic optical element (HOE) to provide a Maxwellian view, which converges the coherent parallel rays at the center of the human eye and projects the rays directly on the retina. Direct retinal projection does not rely on the eye's crystalline lens, so the virtual object can be seen at any convergence. Conventional binocular stereoscopic displays

(head-mounted displays, for example) disassociate accommodation and convergence, because they have a fixed focal length. This project's stereoscopic retinal projection overcomes that defect. In addition, the HOE grating that provides the holography effect is very efficient in both transparency and diffraction, so users can see the real world and virtual objects at the same time.



Focus to real world



Defocus to real world

Collaborators

Eiji Shimizu
Osaka City University
Laboratories of Image Information
Science and Technology

Hideya Takahashi
Osaka City University

Susumu Bandou
Osaka City University



Display and observer

RV-Border Guards: A Multi-Player Entertainment in Mixed-Reality Space

Toshikazu Ohshima
 Mixed Reality Systems Laboratory, Inc.
 ohshima@mr-system.co.jp

RV-Border Guards is a collaborative shooting game based on state-of-the-art mixed-reality (MR) technology, which seamlessly merges real and virtual worlds. Three players wearing see-through head-mounted displays earn points by shooting virtual invaders. They are armed with virtual gear (helmets and guns) and can intuitively interact with the MR space using gestures. The system is also designed so that the audience

can experience and enjoy the MR entertainment by viewing the entire game field on large monitors. This demonstration of MR technology surpasses conventional virtual reality technology and illustrates a quantum leap in the field of entertainment.

www.mr-system.com/rvbg2000



Physical game field



Bird's eye view of the RV-Border Guards space




RV-Border Guards card



Player's view

Collaborators
 Kiyohide Satoh
 Hiroyuki Yamamoto
 Hideyuki Tamura
 Mixed Reality Systems
 Laboratory, Inc.



Virtual creatures

VaRionettes is an interactive installation that demonstrates simple and natural control of complex actions and gestural expressions of a virtual actor (an avatar) using a vision-based interface. Users control a virtual actor directly with their hands, without the use of markers, gloves, sensors, or wires. Images of users' hands are captured by a pair of video cameras and processed to recognize gestures and extract 3D pose and motion parameters. This visual information is translated into commands and parameters that control the avatar's posture, actions, and behaviors. Users can control the positions, motions, and gestures

of the avatar's hands, its head direction, and its walking or running behaviors. The avatar's responses are instantaneous and precise, due to the speed and accuracy of the Bell Labs Gesture Recognition technology. This system provides a simple and intuitive interface to a rich set of behaviors that enable complex interactions with other virtual actors in a virtual theater or with other participants in a virtual conference.



Dancing Avatars

Collaborator
Michael Potmesil
Lucent Bell Labs

V-TOYS is a visually interactive robot that understands and reacts to human presence and visual messages. It has a rotating neck, two controllable eyes, controllable, deformable eyebrows, and a mouth. Its video camera recovers visual information and searches for human presence. When it locates a person, it continuously analyzes facial

deformations to find clues about the human's mood. These basic capabilities are augmented by a set of simple behaviors designed to attract a viewer. V-TOY deforms its facial features as it talks, to invite the user to look at it. It mimics, exaggerates, and smiles at user facial deformations, which generates human curiosity.




Figure 1: The V-TOY robot mimicking different facial expressions.




Figure 2: Interactions between a child and V-TOY.


Collaborators
Myron Flickner
Dave Koons
Alex Cozzi
Ismail Haritaoglu
IBM Almaden Research Center

X'tal Vision (Crystal Vision) is a projection-based augmented-reality system composed of a projector with a small iris and a retroreflective screen. It allows users to observe stereoscopic images with almost correct occlusion relationships between the virtual and the real environment. The small iris mounted on the projector has an extensive depth of field, which enables the free arrangement and free shape of

the screen. X'tal Vision was demonstrated at SIGGRAPH 98 and SIGGRAPH 99. By extending the technology to telecommunication, X'tal Head allows users to observe a stereoscopic head image of a remote person. A head-tracked camera with a constant-orientation link captures the head image, and the system displays a novel, realistic "talking head."




A constant-orientation camera with an operator.



A cubic retroreflective screen, which is attached to a five-degree-of-freedom manipulator.



Projected image. Both the image and the cubic screen are controlled to follow the remote person's head motion.



Collaborators
Naoki Kawakami
Dairoku Sekiguchi
Yasuyuki Yanagida
Taro Maeda
Susumu Tachi
The University of Tokyo




A prototype of a constant-orientation camera.


This demonstration of automatic identification and data capture technologies and applications is essentially a “people tracker.” It captures location information on individuals as they move through SIGGRAPH 2000 Emerging Technologies.

As they leave the program space, users can view their individual paths, information on other people who are now or have been in the space, and general traffic-pattern information for Emerging Technologies. Interactive viewing is available on a kiosk similar to a “You Are Here” map, except that it traces individual paths, so it becomes a “You Were There” map.


Visualization of the overall traffic pattern data and of an individual person’s movement are accomplished through both 2D and 3D graphic applications that analyze volume and flow changes over time, patterns of specific demographic groups, and more. In this way, the technology demonstrates useful applications for traffic-flow planning and real-time traffic adjustments, but it also raises some key privacy issues, because information is captured without the users’ knowledge or permission.



Sample individual track through mock exhibit space.



Visualization of time-accumulated sample densities.




Volumetric visualization of density data.

Collaborators

Paul Beardsley
Hanspeter Pfister
Alyn Rockwood
Mitsubishi Electric Research Lab

Janet McAndless
Capstone Systems Inc.

Steve Van Frank
Van Frank Consulting



Contents

- | | | | |
|-----|--|-----|---|
| 106 | The Actual Reality of Virtual Reality Software | 138 | No Art Jargon! |
| 110 | Computer and Video Games: The Next Generation | 140 | Rich Media for the Web |
| 112 | Designing for Convergence in the Entertainment Industry | 142 | Understanding the Past: Computer Graphics and Archaeology |
| 115 | Digital Cel Animation in Japan | 145 | Visualizing the Cosmos: Smoke or Mirrors |
| 118 | Emotional Simulator: The Tears and Fears of Creating a Compelling Simulated Experience for Both Entertainment and Training | | |
| 122 | Exploring New Roles for Interactive Virtual Characters | | |
| 125 | The Healing Powers of Virtual Reality | | |
| 128 | How Free is Free Software? | | |
| 131 | Interactive Storytelling: New Genres & Directions | | |
| 135 | James Brown: Putting a New Face on the Godfather of Soul | | |

SPECIAL SESSIONS

- | | |
|-----|-------------------------------------|
| 148 | Art & Culture Papers |
| 150 | Fiction 2001 |
| 153 | The Making of "The Perfect Storm" |
| 155 | Phil Tippett's History of Animation |
| 156 | Web3D RoundUP |

PANELS

ALIZA CORSON
WALT DISNEY FEATURE ANIMATION



Committee

CHAIR

Aliza Corson
Walt Disney Feature Animation

JURY

Wayne Carlson
The Ohio State University

Isaac Kerlow
The Walt Disney Company

Claudio Pinhanez
IBM T.J. Watson Research Center

Kathryn Saunders
Royal Ontario Museum

Bill Sherman
National Center for Supercomputing Applications

Joan Staveley
Faust Logic Inc.

PROGRAM COORDINATOR

Vicki Caulfield

SPECIAL SESSIONS COORDINATOR

Betsy Asher Hall
Rhythm & Hues Studios



Introduction

TODAY'S ISSUES ARE LARGE AND MULTIDIMENSIONAL.

Researchers and artists from all parts of the world, whose fields of specialty are as diverse as their cultural backgrounds, are collaborating in ways never before possible via the Internet. Questions emerge regarding standards, access, and truthful representation of data. Experts contemplate the future implications of what has been set in motion today.

The presenters of SIGGRAPH 2000 Panels and Special Sessions tackle these issues in art, animation, archaeology, gaming, engineering, and medicine. The focus begins to shift from one person using a computer to the relationship between two or more people using the computer as a conduit. Virtual environment research explores territory beyond the horizon, yet the applications are widely varied from military training to interactive storytelling; medical diagnosis and treatment to the exploration of ancient cities.

In this section of Conference Abstracts & Applications, panelists summarize their positions and speculations on the future of computer graphics and interactive techniques.

Thanks to the organizers, moderators, and panelists, for sharing their thoughts and brilliant work with us. Thanks also to an exceptional Panels Jury for their guidance in bringing together this enlightening and provocative program.



In the early 1990s, the graphics community widely acknowledged the promise and potential of large-scale immersive environments. We envisioned myriad opportunities to use VR to transform the ways in which we visualize and interact with complex information. Two key years, 1992 and 1993, clearly marked the acceptance of immersive environments, with the introduction and embrace of the Virtual Portal², the CAVE¹, and the Responsive Workbench³. These systems helped establish the field of virtual reality by introducing the novel yet pragmatic use of proven, familiar projection systems to create immersive displays.

However, immersive virtual reality systems comprise not only a novel display environment. They also demand a new paradigm for interaction with information and with our colleagues and audiences in research, production, and entertainment situations. Since 1995, the number of immersive display facilities has grown at a fast pace. In 1999 alone, some 200 new installations were set up worldwide, in academic, research, industrial, and public sites. Applications include engineering design, medical training, architectural walkthroughs, military simulations, entertainment rides, and remote meetings.

We have yet to reach VR nirvana, however. The organization that builds an immersive display facility faces a significant challenge. While the hardware technologies have evolved, and the business infrastructure has emerged, there is not a single integrated set of commercial or research software tools to effectively and elegantly build and run virtual environments. The consequence? Unless they hire a full-time wizard or two with interface design skills and programming talents, businesses cannot rely on immersive environments for routine use, whether as support tools for scientific and engineering problem-solving, for museum exhibition, or for commercial marketing activities.

Large-scale immersive environments integrate varied hardware and low-level software technologies that simply are not hidden from application developers in today's development environments. Is this lack of transparency a good thing?

Today, every VR facility relies on its own set of customized tools, typically home-grown and incompatible with any other set of tools⁴. This creates duplication of effort, frustrations on the user's end, and delays in accepting new technologies, since it means significant code modifications to existing tools.

On the other hand, there is a growing belief that sophisticated software tools are available that can do "everything" to control a virtual environment, given the proper programming treatment. But each organization has a different set of demands, and no general-purpose toolkit could possibly satisfy those demands. This viewpoint dichotomy is causing much confusion and frustration among a large number of developers who are not interested in dealing with software design and development, but simply want to produce and use immersive environments.

This panel consists of top experts in development and application of virtual environments with very different positions on the current state of software technology.

Bernd Froehlich

"At GMD, we are developing Avango, a software framework for local and distributed VR applications. Our framework has been extensively used for developing large-scale visualization and story telling applications for different virtual environment display systems. While the visualization applications are highly interactive and deal with huge datasets, most of our entertainment applications are extremely complex, involving lots of animated textures and complex models. The development in both application domains relies heavily on a scripting language fully integrated with our Avango virtual environment framework.

"All high-level Avango objects can be created and manipulated from the scripting level. The Avango scripting interface allows a two-track approach to application development. Complex and performance-critical functionality is implemented in C++ by subclassing and extending existing Avango classes. The application itself is then just a collection of scripts that instantiate the desired Avango objects, call methods on them, set their field values, and define relationships between them. Scripts can be written, tested, debugged, and modified while the application is running. This leads to very short turnaround times during the development cycle and provides a very powerful environment for rapid application prototyping. Beginners and occasional users of the system mostly stick with implementing scripts. They don't need to deal with C++ code at all, but they are still able to develop complex scenarios.

“Much of the application development effort has to be devoted to the user interface part. Many virtual environment applications try to map interaction tasks to a single pointing device with a few buttons (equivalent to the omnipresent mouse in 2D systems). In real life, carpenters don’t try to perform all their tasks with a single hammer. Instead, a variety of different tools with special properties are used, each perfectly suited for only a few tasks. In virtual environments, passive real-world props are beginning to take up the role of these special tools. One such example is the Cubic Mouse, an interface device developed for navigation and slicing of volumetric data sets. Application developers need to control, configure, and integrate a variety of these devices, ranging from simple pointing devices to pen-based computers and fully custom props. Virtual environment systems have to provide interaction frameworks to support a mix and match of these very different interaction concepts.”

Bernd Froehlich is the coordinator for scientific visualization projects at the German National Research Center for Information Technology (GMD). His recent work focuses on visualization and interaction techniques for local and distributed virtual environments. He is also heading the VRGeo consortium activities at GMD. The consortium brings together oil and gas companies with developers of software for oil and gas applications. This effort involves development of prototype VR applications specific to oil and gas activities, which enable consortia members to evaluate how working in virtual environments can most effectively be used in the oil and gas industry. From 1995 to 1997, he worked as a research associate with Pat Hanrahan and the Computer Graphics Laboratory at Stanford University, where his work focused on interaction techniques and support for tight collaboration in virtual environments. From 1993 to 1995, he headed the Responsive Workbench project at GMD. He received his masters in science and doctor of philosophy in computer science from the Technical University of Braunschweig in 1988 and 1992, respectively.

Mary Cole

“Virtual environments have great promise, but their acceptance in the oil industry has been held back by the lack of compelling application software. The general-purpose toolkits on which one can build these compelling domain-specific applications have increased in number and improved greatly, even in the past year. However, there is still much work to be done.

“A general-purpose toolkit must be robust, flexible, extensible, and scalable. It must be robust, so as to not further complicate the development process. It must be flexible in supporting enough choices of development language, scenegraph, etc. to be acceptable to a critical mass of application developers. It must be extensible, so that the application developer can add critical functionality not supplied by the toolkit. Scalable toolkits allow the virtual environment developer to build an application that scales from a personal virtual environment, such as a desktop, or HMD, through benches, up to a large-scale collaborative virtual environment, such as a CAVE.

“Our users want applications that they can use both alone in their offices and with groups in larger collaborative environments. The user-interface paradigm may be different in a personal (for example, desktop) environment than in a large-scale environment, but it must have enough conceptual commonality so that a desktop user is not lost when using the application on the larger environment. The general-purpose toolkit must support a wide variety of user interface and tracking devices, and it must support execution-time selection of devices and user-interface modules. The application developer must take care that the user interface paradigms are appropriate to the scale at which the application is being used and consistent and analogous between scales. We need to develop human interface guidelines that scale consistently across the breadth of different virtual environments.”

Mary Cole is a visualization software developer at Atlantic Richfield Compa. Currently, she is responsible for developing immersive applications for exploration and interpretation of geophysical data. Prior to ARCO, Mary worked at CEMAX, a medical imaging company. She graduated from MIT with a BSCS degree.

Carolina Cruz-Neira

“Applications for virtual environments have been developed since the very early stages of virtual reality technology. In those early stages, most applications were built from scratch, creating a large range of software tools that were very difficult to reuse on other applications. In 1992, as part of my PhD dissertation, I designed and implemented the CAVE libraries, a software library for development of immersive applications that provides developers hardware independence. Today, the CAVE libraries are being used by several thousand users and are the basis for a wide diversity of application domains. These libraries were based on software engineering techniques available at that time, and provided enough support to allow the CAVE and other virtual reality systems to be used in a variety of fields.

References

1. Cruz-Neira, C. et al. (1992). The CAVE audio visual experience: Automatic virtual environment. *Communications of the ACM*, June 1992, 65-72.
2. Deering, M. (1993). Making virtual reality more real: Experience with the virtual portal. *Proceedings of Graphics Interface'93*. May 1993, 219-226.
3. Frshlich, B., Grunst, G., Kruger, W., Wesche, G. (1995). The responsive workbench: A virtual working environment for physicians. *Elevier Science, Computers in Biology and Medicine*, 25 (2.) 301-308.
4. Lastra, A. (1997). Programming virtual worlds. *ACM SIGGRAPH 97 Course 29*.

“Very recently, taking advantage of more advanced software engineering techniques such as object-oriented design, component-based approaches, and scalability, we have made available to the VR community VRJuggler(www.vrjuggler.org), a set of libraries for immersive application development. VRJuggler provides basic functionality and transparent portability across VR platforms, and it allows users to add functionality and expand its capabilities. Although VRJuggler is not the solution for all applications, it can certainly provide a basis to develop a more sophisticated tool.

“To achieve the goal of a common software development environment, some toolkit or library needs to be accepted as the standard and serve as the platform for more sophisticated tools, general or specialized, to be developed. Today’s software for VR offers a few options to choose from for that standard: VRJuggler, Bamboo, and Avango are example of research software that is mature enough to be seriously considered as potential standards. There is good software available for developing VR applications, but it is being limited by other factors. Two of these factors are the lack of agreement in the community on what model could integrate the different software pieces and the commercialization efforts for the individual pieces.”

Carolina Cruz-Neira is Litton assistant professor in the Electrical and Computer Engineering Department and the associate director of the Virtual Reality Applications Center, at Iowa State University. Her main research area is software models to integrate the variety of hardware components needed to develop an immersive system and to provide application portability and scalability across platforms and across collaborative situations. She obtained a PhD in electrical engineering and computer science from the Electronic Visualization Laboratory at the University of Illinois at Chicago in May of 1995. Her PhD research involved design and implementation of the CAVE and development of paradigms to combine high-performance computing and communications with the CAVE. She received her masters degree in EECS at the University of Illinois at Chicago and graduated cum laude in systems engineering from the Universidad Metropolitana in Caracas, Venezuela.

Linda Jacobson

“The virtual reality developers scene works the same way that social and political movements work: Pioneers make sacrifices so the next generation can move on, unencumbered by bitter memories of the early days, free of old demons. We are the pioneers. Therefore, we suffer. Thanks to us, future immersive environmentalists will easily wield powerful, commercial software tools to create cost-effective, compelling models and environments with quantified return on investment.

“On another level, the VR developers scene operates the same way that the business world operates. Pioneers compete, repeat, and occasionally stumble, and only a minority live long enough to outfit the next generation. We, the pioneers, have enabled and endured the evolution of the hardware. The technologies improved because of the feedback we provided to the creators. Now, we’re moving on to the software, facing the facts of its infancy: The agony of data conversion. The schizophrenia of interaction methods. The murky darkness of low-level code. Whether we build the software tools ourselves or use the tools that others create, we must accept our role and responsibility as pioneers. We must use what we’ve got and make it count. We must prove and document the efficacy of the tools and what we do with them, as we admit our work’s weaknesses and strive to eliminate those weaknesses.

“As pioneers, we have an obligation to provide constructive feedback to the source of the tools. In today’s world, if you can help the creator demonstrate and document the tool’s cold, hard value, then the money will come to enhance and improve the tool. And whether we’re talking about socio-political movements or business operations, money is the fertilizer that encourages young things to grow. We have the software; the real question is, where is the money?”

Linda Jacobson is SGI’s virtual reality evangelist and a senior manager for SGI visual solutions marketing. She spearheaded SGI’s Reality Center business activities and continues to develop corporate strategies concerning advanced user interfaces for visual computing. Previously a journalist, author and educator, she is the author of “Garage Virtual Reality” (1994, Sams/Macmillan) and “CyberArts: Exploring Art & Technology” (1992, Miller-Freeman). In addition, she helped create and launch several magazines, including *Wired* and *Computer Life*, and in 1993 co-founded the Virtual Reality Education Foundation. She has been a faculty member of San Francisco State University’s multimedia studies program, a columnist for CNET, a talk-show host on the nationally syndicated radio program “On Computers,” and a patent application writer for Xerox PARC. She has also worked as a digital puppeteer, using performance animation systems in theater and music productions.

Ken Pimentel

“The lack of robust, intuitive software for immersive environments greatly limits who can use these environments and the problems they address. Immersive software is written without any common style guides, radically different interaction paradigms, and no commonality for how primitive elements (such as buttons, or dials) are displayed. Every time new users step into an immersive environment, they must be taught how to navigate, how to interact with the interface device, how to manipulate the scene, and dozens of other varying tasks.

“This presents a staggering cognitive load, ensures that only application experts can derive the value of the visualization experience, and that these experts can’t use their knowledge on the next immersive application. It is much like walking up to a computer and within a single session having to use Windows, the Mac OS, DOS, Unix, and sometimes a video-game interface, and then, for each environment, you have to replace your mouse with some other interface device. This presents an impossible barrier to widespread use and effective interaction within immersive environments.”

Kenneth Pimentel is director of immersive solutions for Engineering Animation Inc. Currently, he is responsible for understanding the needs of the manufacturing industry and applying appropriate immersive techniques to those needs. Prior to EAI, he was vice president of business development for Sense8 and co-author of the book “Virtual Reality: Through the New Looking Glass.” Prior to Sense8, he worked at Intel for 10 years. He graduated from the University of California, Davis with a BSEE degree.

Kent Watsen

“Development of virtual environment (VE) applications is a time-consuming effort understood by few. While the lifetime of some applications can be measured in days or weeks, a greater return on investment is obtained if the software can be repurposed over the course of months and years. Recognizing this, many general-purpose toolkits have been produced to facilitate the realization of both application-independent and domain-specific abstractions. While these toolkits have indeed accelerated the development life-cycle for many individual efforts, they have also implicitly slowed down the VE community as a whole. Because these toolkits have different architectural implementations, they have fragmented the VE community. The result: It is not possible to re-use code (source or compiled) across their boundaries.

“This scenario is identical to that in any market that lacks standards for interoperability. If the VE community is ever to exceed our current growth rate, in terms of complexity management, this issue will have to be resolved. The solution to this dilemma is not technical, but political. We must set aside our not-invented-here egos and business plans and realize that there is a greater benefit for all if we collectively move to a razor-and-blades model. That is, in business terms, we should give away, adopt, or share a common underlying platform (the razor) on which we can then specialize via independently developed (and potentially licensed) extensions (the blades).

“Many will perceive this idealistic perspective as theoretically interesting but impossible to achieve in practice, as it is hard enough to agree upon things as fundamental as the programming language or the scenegraph to use. However, component-oriented programming, a recent trend in software engineering focused on establishing the ‘standards for interoperability’, may be the proverbial ‘silver bullet’ (Brooks, 1987) the VE community so desperately needs. The idealistic solution may well be achievable today if the shared underlying architecture (the razor) is a component framework and the independently developed extensions (the blades) are components.”

Kent Watsen is interested in all aspects of making social virtual environments an integrated part of future generations. To this end, he is concerned with all of the technical, social, and psychological issues that are the mainstay of the relevant conference and publication venues. His experience is the result of almost a decade of academic and commercial research. Specifically, his undergraduate thesis, “Raytracing for Realtime Animation,” was the primary reason why DCS Corporation hired him out of school to develop VisualWorld, a performer-based, DIS-compliant visual simulation toolkit. After developing this toolkit, he co-developed EasyScene for Corypheus Software (a.k.a. Centric Software), while also designing and developing both the “OpenSeas” and the “EasyCharacter” plugin modules. He has since moved on to pursue his PhD under Mike Zyda at the Naval Postgraduate School while acting as project manager of the NPSNET Research Group in the computer science department. His dissertation topic, “Bamboo,” is an open source, platform- and language- independent component framework on which VE-specific components are being developed.

In 1999, computer and video game revenue topped \$7.4 billion dollars in the US (a 20-percent increase over 1998). This industry is set to eclipse the feature-film market, which had US revenue of \$7.5 million in box office receipts, which were boosted by increased ticket prices (hardware and software prices have declined on average).

Advances in computer technology have fueled a rise in the quality of computer and video games, and they can now display cinema-quality graphics. Greater player control is on the horizon, and emotional connections between player and game will soon surpass those in feature films.

In this panel, we examine the history and present state of game development, then focus on the future of the industry to see how increased computing power, metadesign, media convergence, and online access will change the face of this industry. Finally, we examine the moral obligations that developers face when they determine content.

NEW DEVELOPMENTS

Development of computer and video games has always been closely tied to computing power. As power and capacity skyrocket, so do game capabilities and audience expectations. Sony leads the movement to real-time interaction with its release of PlayStation 2, but the fierce competition for a piece of the \$16 billion worldwide revenue pie ensures that its leadership will be challenged within the year, if not next month.

VISIONS OF THE FUTURE: SUPER DEVELOPERS

The promise of cinema-quality graphics has already led to important predictions on the evolution of super developers. This prediction is based on:

- A need for high-end computer graphics modeling capabilities, large poly-count/NURBS models, and multi-layered textures.
- Real-time interactive lighting and procedural particle animation.
- Sophisticated personality design and performance with greater detail of movement and expression.
- Cinematic-quality storylines comparable to film and TV, and access to their properties.
- Freeform movement as opposed to pre-calculated poses and movement cycles.
- Increased development costs and time, which means that ventures will require greater risk and capitalization than is available to most developers.

EMOTION SYNTHESIS

A character's reaction to the environment and unique developments in the story line coupled with the player's choices and abilities will create a seemingly emotional play. The dream that computers will mimic human thought and emotion seems like a reasonable goal. Simultaneously, multi-player games will inject a great deal of human quality to these games as players control their game counterparts. Add to this an audio connection that enhances communications between players, and the experience becomes multi-layered.

MEDIA CONVERGENCE

As signaled by the merger of Time Warner and America Online, the success of a game will depend on its presence in a variety of media, including print, film, television, and online. Each of these media will contribute a unique perspective that provides greater understanding of the story and characters in a game instead of just duplicating one predominant medium.

ONLINE CAPABILITIES

Online connections will allow quicker access to game updates and the possibility of interaction between players from opposite corners of the world. The Web is potentially the lifeline between consumers and game developers. The possibility that developers may be able to circumvent distributors by way of the Internet also begs the question: Is there room for small developers? Will the relationship between major film studios and independents provide a model for the game industry?

PROGRAMMING

On the horizon: potentially radical changes in the way we approach computing. New methods will supplant the need for endless strings of code and predetermining every aspect of a game. Two obvious solutions are:

- Increased encapsulation for smaller module changes with the most dramatic alterations to the entire structure of the program.
- Rule-oriented programming that replaces vector graphics and its hunger for storage space and computing power.

MORALS AND ETHICS

All forms of entertainment will be scrutinized for their violent and sexual content, just as television and film are. A ratings system is in place, but as games become more predominant, they will absorb a greater amount of criticism. What responsibilities do developers and publishers have to the general public? What can they learn from other forms of media and their past battles?

Panelists

Lucy Bradshaw
Electronic Arts

Tom Hershey
Sony Pictures Imageworks

Henry Jenkins
Massachusetts Institute
of Technology

Dominic Mallinson
Sony Computer
Entertainment of America

Christopher Weaver
ZeniMax Media Inc.

Lucy Bradshaw



Lucy Bradshaw is general manager of Electronic Arts' development studio Maxis, producers of The Sims. Prior to this position, she served as executive producer of the studio's SimCity franchise, overseeing the development of SimCity 3000, 1999's top-grossing PC game, and SimCity.com. She got her start at Activision, then moved to LucasArts to be associate director of development. While there, she worked on

"Monkey Island II," "Indiana Jones and the Fate of Atlantis," "Star Wars Nintendo," and "Secret Weapons of the Luftwaffe." She is a graduate of the University of Michigan, where she earned a bachelor of arts in psychology and has two children.

Tom Hershey



"Interactive storytelling has lacked the effective pacing of a good movie or book, and interaction with technology-limited characters left players with a less engaging experience. The new technologies now available to us will change all that. Compelling story and characters in a game will be important, but the ability of the player to create memorable and unique experiences

will put certain games over the top. New technology will radically improve the look of games and allow for more sophisticated interaction of player, character, and environment. The less authors have to 'spoon feed' the players, the greater the playability and enjoyment."

A graduate of MIT, Tom Hershey worked as a programmer, specializing in graphics applications for PCs. Recruited from the UCLA MBA program in 1988, he joined Columbia Pictures and worked for four years as director of production administration, helping to oversee production of feature films. He now leads Imageworks' movement into content development for Sony's PlayStation 2.

Henry Jenkins



Henry Jenkins is the Ann Friedlaender Professor of Humanities and director of Comparative Media Studies at the Massachusetts Institute of Technology. The central focus of his research is a study of children and adolescents in a mediated culture; his most recent books include "From Barbie To Mortal Kombat: Gender And Computer Games" and "The Children's Culture Reader."

He holds a PhD in communication arts from the University of Wisconsin-Madison.

Dominic Mallinson



"Within the next year, we will have cheap processing power to manipulate complex 3D geometry and animations of fantastic quality, a multitude of digital consumer electronics from digital cameras and camcorders to MP3 players, and widespread broadband networking in the home. The big question is: Will we have software content of equal quality?"

"There are key technologies that will help create this content. Advanced rendering, physical simulation/kinematics, artificial intelligence (behaviors), and digital IO will be essential to believable interaction since it will no longer be acceptable for games to have characters with pre-scripted animations and behaviors. Artificial intelligence should be used to prevent repetition. Characters should change their behaviors according to the state of the world and the player's interaction."

After graduating in computer science from the University of Durham, Dominic Mallinson spent time at Microsoft working on their first C++ compiler. He returned to the United Kingdom to work for Pilkington Glass on CAD and factory automation, and eventually, he joined Psygnosis and remained there for seven years, during which time Sony acquired the company and eventually launched Sony PlayStation.

George Suhayda



A graduate of Clemson University and the Yale School of Drama, George Suhayda joined Sony Imageworks in 1998 and has worked on "Contact," "Snow Falling On Cedars," "Sphere," "City Of Angels," and "What Planet Are You From?" Currently, he is helping Sony Imageworks develop games for Sony's PlayStation 2.

Christopher Weaver



"The computer games industry has undergone tectonic shifts over the past 25 years. Now, with the advent of even more powerful graphics cards and the new breed of consoles pushing millions of triangles per second, the future will be limited by imagination, not graphics. But do such hardware enhancements necessarily make for more successful games?"

Founder and former president of Bethesda Softworks, Christopher Weaver also participated in some of the earliest development of interactive multimedia. He was chief engineer for the Congressional Subcommittee on Communications, vice president of science & technology for the National Cable Television Association, and director of the Office of Technology Forecasting for the American Broadcasting Company. He holds graduate degrees from Wesleyan University and MIT. He was appointed a fellow of the Robotics Simulation Laboratory at Carnegie Mellon University in 1995.

The digital revolution has permeated virtually every aspect of modern society, and today's Internet has exploded far beyond its original scope to irrevocably alter the landscape of contemporary media and entertainment.

Computing, television, and communications technologies continue to rapidly converge and have enabled some interesting hybrids. But as fast broadband connectivity pervades the Internet, truly unique and powerful creative synergies are emerging, with dramatic implications for digital entertainment.

The synthesis of such high-speed, always-on digital technologies is collectively referred to as "convergence." But convergence implies more than just the gee-whiz factor of its technological underpinnings; indeed, it is poised to become *the* predominant cultural and lifestyle issue of our contemporary wired society.

Consequently, creative media professionals face myriad challenges — and opportunities — as they begin exploring what it means to "design for convergence." And companies, from traditional conglomerates to nascent Web startups, find themselves overlapping and infiltrating each other's disciplines as they seek to define experiences that successfully merge entertainment, participation, and ubiquity.

Panelists from a variety of multidisciplinary backgrounds and perspectives discuss conceptual methodologies and illustrate them with practical examples of current broadband experiments and works-in-progress, exploring such questions as:

- What does it mean to design rich content and useful context within this budding, evolving environment? What do consumers want, and what are the cultural implications?
- How must companies adapt and extrapolate their existing strengths, and how must individual skill sets mutate to enable design and implementation of future systems?

- What types of content and experiences might take most appropriate advantage of these new synergies?
- And ultimately, what might be the "killer app" of convergence?

CONVERGENCE CONTENT STRATEGIES

Michael Gough

"Many or most of the strategies for creating content for emerging media platforms are either device-driven (innovative new uses of this technology or that platform) or driven by the imperatives of an existing medium. Device-driven media producers are enamored with the technology or the technique. They create things because it is possible to create them. Media producers who are endeavoring to port content from its parent medium to the new platform(s) are often stuck with the priorities as well as the prejudices of the medium they spring from. A more radical (and appropriate) approach is informed by understanding that emerging technology is in part a response to its potential consumer. In this view, technology's creators (that would be all of us) remain in power, choosing to create because of changing needs or wants or wishes.

"What follows logically is a constantly evolving creator/consumer who, if understood, can help provide a road map for emerging media. By focusing our efforts to create convergence content on an understanding of the emerging media consumer, we can successfully produce experiences that will be consumed.

"As a result of this understanding, there are a few assumptions that much of our work is based on. These are rooted in 'cultural memes' that provide insight into the nature of the emerging media consumer. We have identified four such memes:

1. Multiple simultaneous input (replaces mono or modality (channels).
2. Technology lust (replaces technophobia).
3. Participation necessity (replaces couch-potato presentation).
4. Virtual presence (simulacra) (replaces living-room community).



"James Cameron's Titanic Explorer," an interactive documentary.
Rare Medium, for Fox Interactive and 20th Century Fox.

Panelists

Michael Gough
Quokka Sports

Don Guy
Don Guy & Associates

Jesse Schell
Walt Disney Imagineering

Kevin Townsend
Rare Science + Fiction

“The next (convergence) medium will be created in response to these memes.”

As chief creative officer, Michael Gough is the force behind Quokka’s immersive sports experiences, leveraging data, images, video, and next-generation presentation methodologies to collapse the distance between the event and the viewer, breaking down the barriers between vicarious experience and the feeling of actual participation. He joined Quokka in July 1997 as vice president, design and creative director and was promoted to chief creative officer and executive producer in September 1998. In August 1995, he co-founded Construct Internet Design, a digital media design firm, where he served as creative director until July 1997. Prior to that, he co-founded Jones, Partners: Architecture, a design-focused architecture firm, where he served as managing partner from December 1994 to August 1995. Earlier in his career, he was an architect for Holt Hinshaw Pfau Jones and, before that, an architect for the San Jose Redevelopment Agency. He studied architecture at California Polytechnic State University.

Don Guy

“Directors working in ‘traditional’ entertainment media such as film and television have inherited a rich tradition of visual and narrative storytelling techniques. Computer-generated imagery and digital image manipulation have enabled a quantum leap forward in the possibilities of these techniques, and as a result have profoundly affected the next generation of storytellers.

“As audiences became more visually sophisticated and media-savvy, and as the simplification and democratization of digital technologies continues to penetrate the consumer sector, children come into contact with them. Our children in turn have become some of the earliest adopters of the available technology, and this has nurtured and expanded their storytelling aspirations.

“Children are unaware of television and the Internet converging into one interactive experience. They are unaware of the premise or promise of high-speed broadband networks. However, because

children have been raised in this climate of a digitally interactive language, there has been a profound impact on the development of their visual and narrative methodologies. Today’s wired youngsters will reap the full benefits of a mature digital infrastructure by exercising their own voices through the latest technologies. A survey and analysis of storytelling from children’s perspectives — and their visions of the future — can provide a valuable window of insight for today’s convergent content and technology developers.

“How are tomorrow’s media practitioners capitalizing on this enormous potential for self-expression?”

Don Guy is a cinematographer with 20 years of experience in not only film and commercial production, but also creative direction, marketing, and executive management. His clients have included ITT, General Motors, Ford, IBM, Coca-Cola, the Smithsonian Institution, and the American Dairy Association. His work has received numerous Clio and Cannes Lion awards, and he also received an Academy Award nomination for Best Documentary Short. His film “Portrait of a Railroad,” for Burlington Northern, won 17 international film festival awards. He has a strong interest in digital technologies, including motion-control photography, digital film techniques, and computer-generated imagery. He is an MFA candidate at the UCLA School of Theatre, Film & Television, and has studied with Lee Strasberg’s Actor’s Studio Director’s Workshop in New York. He received his BA in liberal arts from Goddard College.

Jesse Schell

“The convergence of cheap computing power and high-bandwidth networking to the home will change interactive entertainment from a niche-market novelty to a mass-market medium. This convergence will have profound effects on the entire entertainment industry, and on the personal lives of each and every one of us. Insights into human psychology are the best path to making successful predictions about what will succeed and what will fail in this new world of connected entertainment.



Quokka Sports’enhanced HDTV coverage of Olympic men’s cycling.

“Disney’s VR Studio has created multiplayer virtual reality attractions such as ‘Aladdin’s Magic Carpet Ride,’ ‘Hercules in the Underworld,’ ‘Mickey’s Toontown Tag,’ and the soon-to-be-released ‘Pirates of the Caribbean.’ This presentation analyzes aspects of group psychology that made these attractions successful and presents cautionary examples about multiplayer games that failed to take psychological factors into account.”

Jesse Schell is a show programmer and game designer at the Walt Disney Imagineering Virtual Reality Studio. He has helped develop such attractions as: “Aladdin’s Magic Carpet Ride,” a head-mounted, display-based VR attraction currently installed at DisneyQuest; “Hercules in the Underworld,” a CAVE-based VR attraction, also at DisneyQuest; “Mickey’s Toontown Tag,” a multiplayer game currently installed at Epcot Center; and the interactive “Pirates of the Caribbean” multiplayer, CAVE-based VR attraction, soon to open at DisneyQuest. He has appeared in several panels and courses, including: Film Festival Futura 00; The Future of Interactive Entertainment, SIGGRAPH 99; Fiction 2000: Technology, Tradition, and the Essence of Story; CGDC 99; and SIGGRAPH 98, Immersive Environments: Research, Applications, and Magic. He has a BS in computer science from Rensselaer and an MS in information networking from Carnegie Mellon. His main interest is in making virtual worlds more *fun*.

Kevin Townsend

“A new medium will have its stars, shows, and studios.”

“In the next few years, a truly user-controlled, rich media experience will be realized, an experience that delivers top-quality entertainment within a robust, sophisticated software environment. These programming-and-services-on-demand experiences will finally deliver their long-projected and proportionately huge impact. The creative and business forces associated with these experiences promise to dismantle and rebuild the business and content models that have been groomed over the last 50 years in the media and entertainment arenas.

“This new set of forces will give the audience programming and services they want and need, when they want them, where they want them, and on all the devices that occupy their lives. Companies that can predict and respond to these forces are being built right now. They will produce a new class of content and services showcasing existing and new talent. They will look like game developers and Web agencies and live-action production companies and broadcast design studios and visual effects shops, all under one roof.”

Kevin Townsend is a media veteran with 16 years in the entertainment industry, including work in film, television, commercials, and music videos. He spent five years with Industrial Light + Magic, where he headed both the commercial and new media divisions. In this capacity, he was also involved with various professional presentations of ILM’s work, including those at SIGGRAPH conferences. Currently, he builds companies and divisions of companies, specializing in combining technology with entertainment.



Disney’s “Pirates of the Caribbean,” a multiplayer CAVE-based VR attraction.



Early adopters: nine-year old children creating digital postcards for a school report of their trip to the Mayan ruins at Chichen Itza. (courtesy Don Guy)

Japanese cel animation has become very popular all over the world. In recent years, extensive use of computer graphics technology in cel-animated production has led to some innovative technical developments and changes in the “traditional” animation production pipeline in Japan. This panel focuses on how computer graphics technology has been used in practice and how it could be used more widely in Japanese entertainment, not solely for new visual effects in animation.

Western applications of digital technology in animation are quite different from those in Japan. These variations are primarily due to differences in history, culture, taste, and style. What are the practical differences in the use of digital technology? How are new digital techniques developed for Japanese animation? This panel answers these questions through case studies featuring Japanese digital cel animation created by leading production companies.

Topics include:

- The history of Japanese traditional cel animation.
- Effective application of current computer graphics technology:
 - Mixing 3D computer graphics with cel animation: “Pokémon” (OLM, Inc.) and “Tekkon Kinkreet” (Trilogy Co.)
 - Non-photorealistic rendering for digital cel animation: “Princess Mononoke” (Studio Ghibli) vs. “Prince of Egypt” (DreamWorks SKG)
- How can animators use digital technology to make Japanese animation enjoyable?
 - Case studies from the animators’ and directors’ points of view: “Princess Mononoke” and “My Neighbors the Yamadas” (Studio Ghibli)
- What are the traditional and current problems faced by Japanese animation and how can they be solved?
 - Case studies of Softimage and Hitachi
- Future directions for computer graphics technology in digital cel animation

Ken Anjo

“Computer graphics researchers find a variety of very interesting, exciting, and unexpected research topics in the field of digital cel animation. In particular, when we consider Japanese cel animation, we notice the uniqueness and richness of Japanese entertainment, which could be much more fully explored with the help of digital technologies.

“Animators who use digital technologies should consider two basic requirements: reducing costs and creating new effects. These two requirements are usually a trade-off, and, for most animation production houses in Japan, unlike Studio Ghibli and Trilogy, the first step in making digital cel animation would be to fully use the currently available commercial software, but to do in a traditional way, depending on their hand-drawing skills. Digital software products are valid only when animators create visual effects that cannot be made without the software technologies. But what is the actual situation in Japan? Especially for production managers, creating a cost-effective digital cel animation system is a pressing need, so they must develop custom methodologies that are well suited for Japanese animation.

“In this panel, I show a few case studies taken from the recent works of OLM, Inc. The case studies highlight comparison of two Pokémon movies: “Mewto Strikes Back!” and “Pokémon, The First Movie.” The first is Japanese cel animation, with a few digital effects, while the second includes about 80 digitally processed shots. Actually, the second movie is a remake of the first one, with new stories and many digital effects, so I demonstrate the different nuances among the two movies, which are the result of digital effects. The digital techniques used are relatively basic and not so tricky that they cannot make innovative effects. However, it is very interesting to show how the Japanese cel animation is translated into the digital form.”

Ken Anjo is the technical supervisor at OLM Digital, Inc. He has also been a guest professor of Keio University since autumn 1999. Prior to joining OLM Digital in 2000, he was a research scientist at Hitachi, Ltd., where he led R&D teams in computer graphics and published many scientific articles, including three SIGGRAPH papers.



Figure 1. Frame from “My Neighbors The Yamadas.”

Michael Arias

“Without a doubt, the rapidly-growing presence of 2D and 3D computer graphics in the cel animation industry workplace, particularly in Japan, has required cel-animation production houses and professionals to adapt well-established production methods to take advantage of digital tools. In addition to various technical obstacles and benefits, though, the new software and hardware systems have also introduced some specialized financial burdens (familiar to any computer graphics producer) and, most interestingly, social challenges related to the merger of computer graphics and traditional animation cultures.

“Having witnessed first-hand the transition from photomechanical-photochemical means to the digital techniques that swept the visual effects industry several years ago, and having been involved in a variety of both Japanese and North American cel-animation productions, I am in a unique position to comment on the impact this technology is having on the cel-animation industry in Japan.

“Sweatshop-like cel animation houses proliferated during Japan’s post-war years as the film and television industries sought cheaper means of content production (which they achieved, largely through production techniques established by director Osamu Tezuka, much to the detriment of the animators themselves). Though the sweatshop days are over, the cel-animation industry in Japan is still plagued by poor working conditions and low animator salaries (relative to the North American and European industries). Payment per-cel is still standard in the Japanese cel-animation industry, but the increasing presence of digital personnel and technologies have begun to force positive changes.

“I discuss my work with both CG and cel-animation staff on the “Tekkon Kinkreet” Pilot, with emphasis on our use of traditional techniques in a CG-dominated production environment.”

Michael Arias began his career in visual effects 11 years ago on the stages of Dream Quest Images as a motion-control camera operator on such films as “The Abyss” and “Total Recall.” Since then, he has worked with directors Douglas Trumbull, David Cronenberg, Spike Lee, and others in the US and Japan. For the past four years, he has been employed as a software developer and consultant by Softimage Co. He lives in Tokyo where he is currently supervising computer graphics at Trilogy Co. for a fully digital animated feature film.

Youichi Horry

“I describe the new vector-based 2D animation system for professional use, which was developed jointly by Hitachi, Ltd. and Japanese animation companies. One of the main features of the system is that it provides solutions to the traditional, hard-to-solve problems in making 2D models from hand-drawing input and in automatic interpolation between the key frames. I also demonstrate our new modeling tool in the system, called magnet modeling. Magnet modeling allows creation of 2D models by Bézier curve fitting. The user simply puts the control points on the input hand-drawn image. Boolean operations for modeling complex 2D shapes are also available in the system. Automated in-betweening is performed with several operations that the user can select, such as pseudo-3D rotation and skeletal displacement. These processes and the system’s GUI are shown along with many applied results, which illustrate the efficiency and novelty of the system.”



Figure 2. Images for “Mewtwo Strikes Back” (top) and “Pokémon, The First Movie” (bottom).

Youichi Horry has been a research scientist at the Central Research Laboratory, Hitachi, Ltd since 1990. From January 1997 to April 1998, he was a visiting researcher with Project SYNTIM, INRIA, Roquencourt, France. His current research interests include hybrid approaches to 2D and 3D computer graphics and image-based rendering and modeling, such as “Tour Into the Picture” (SIGGRAPH 97 and SIGGRAPH 98).

Yoshiyuki Momose

“Though companies and organizations each have their own, different views on digital technology, the transition to digital production in the Japanese animation industry has rapidly progressed in the past two years. The benefits made possible by recent advances in technology have prompted Japanese studios, which have been relying heavily on subcontracting due to chronic labor shortages, to drastically revise production quality and schedule management systems. At the same time, digital technology has stimulated the imaginations of artists who have been, until recently, restricted by the limitations of cel animation.

“The digitization of Studio Ghibli began with the start of the CGI department in 1995, when production began for Director Hayao Miyazaki’s ‘Princess Mononoke’ and has continued through the production of ‘My Neighbors the Yamadas’, Ghibli’s first completely digital feature film. ‘Princess Mononoke’ was the last Ghibli film made with mostly traditional means (only about 10 percent of the scenes were treated digitally). Even digital effects were intentionally subtle, to avoid distracting the audience’s attention. With ‘My Neighbors the

Yamadas’, however, we freed ourselves from the restraints of cel animation. Director Takahata’s feeling from the start was that a strictly cel-based approach would work against the film’s content. In the end, digital technology was the core production tool (not limited to special effects), and it helped realize a visual style that would not be possible otherwise.

“That said, I should stress that I am not interested in pursuing visual novelty with computer graphics. Rather, I am searching for a new style in which we can use CGI naturally as if we were drawing in freehand, with content and technology mutually heightening each other. At this point, it remains to be seen to what extent the introduction of digital technology will encourage new creative efforts, but it is quite clear that digital technology has become something that cannot be ignored by creators from the traditional animation field. I feel this is a good time to go back to the starting point, to rethink and discuss Japanese animation.”

Yoshiyuki Momose is currently the head of the CGI department of Studio Ghibli, where he has a variety of responsibilities (CG creator, animator, and director), depending on the projects at hand. For 25 years, he worked on animated television series and, in the early 1980s, became involved with CGI through the work of Japan’s first CGA production studio, JCGL. In the following 10 years, he worked on various cel-animated projects for Studio Ghibli. Since Studio Ghibli’s CGI department began in 1995, he has worked on CG production for “Princess Mononoke” and “My Neighbors the Yamadas.”



Figure 3. 2D model specified by Bzier fitting (top) and rendered image (bottom).



Figure 4. Comparison of 3D data (wireframe) and final image from “Tekkon Kinkreet.”

Emotional Simulator: The Tears and Fears of Creating a Compelling Simulated Experience for Both Entertainment and Training

Moderator
Christopher Stapleton
University of Central Florida
christopher@stapleton.net

Throw away your thumb-twitching remote and your passive wiggle-box simulator and find out from experts how they see the future of entertainment and research in simulation technology. The diverse industries of entertainment, military training, and scientific research are colliding in order to produce interactive digital experiences that go beyond reality and to what we can only imagine. However, can such diverse cultures speak the same language, share information and collaborate to reach common goals? Are their goals compatible?

This panel brings together experts involved in the intersection of entertainment and simulation research to share their different perspectives on what they see as the obstacles and opportunities in experience-based digital entertainment and the next generation of simulation research.

Representing a “virtual” reality is no longer the challenge in simulation technology. Hollywood has taught us that experiencing compelling simulated scenarios has as much to do with illusion and imagination as it does with reality and scientific modeling. Technology is no longer holding us back in creating the next generation of simulations. We are now leveraging art and science to create compelling and engaging simulations. Artists, scientists, engineers, and warriors need to share the vision with those who are not like ourselves. This convergence is similar to the mindset of Leonardo DaVinci, whose talents spanned many disciplines, including art, science, anatomy, architecture, and the military. In the last century, extreme specialization of disciplines created the greatest creative challenge: interdisciplinary, cross-industry collaboration.

The US Army is now combing storytelling and the art of illusion in compelling realistic scenarios in order to properly train soldiers. Entertainment producers are employing the latest engineering accomplishments to engage an audience in immersive and interactive worlds that allow guests to escape reality. Scientists are becoming artists to successfully immerse collaborators in a shared view of real-time scientific visualization.

As in any new venture, there are difficult paradigms to overcome. Some in the military may feel that a great sensory experience is an exercise in self-indulgence that trivialize the military experience, which involves life-and-death consequences. Scientific research, which depended on the military for its challenges and funding, is now looking to the entertainment industry to find new ways to expand the representational power of computer science. Despite all the simulation technology now available, entertainment has not fully embraced real-time interaction and consequently has limited amusement park attractions to pre-programmed, wiggle-box motion simulators. In the real-time video game market, interaction seems to be stuck in the thumb-twitching paradigm of the TV remote.

Instead of experiencing a renaissance of possibilities in the art of storytelling, we seem to be stuck in repeating yesterday’s blockbusters, and we seem unwilling to risk content that goes beyond thin plots and stereotypical characters. Are audiences ill-prepared for interaction because they’ve been bombarded with passive media? Are creators afraid to invite the audience to join the creative process through interaction?

Our panelists have been challenged with creating the future of simulation technology by integrating scientific simulation research with creative entertainment production to produce simulation technology that entertains, trains, and educates. They are erasing the boundaries between scientific and artistic; practical and playful; real and fantasy applications. These innovators are making the cross-over in order to understand what our commercial, scientific, and artistic communities need to do to prepare for the future.

Christopher Stapleton

Christopher Stapleton is the director of entertainment research at the Institute for Simulation and Training of the University of Central Florida, where he is transferring military simulation technology to education via entertainment. He is also an instructor of interactive entertainment at the University of Central Florida Film Program and the Consortium for



It's a Small World, from the Building Virtual Worlds Class, a part of CMU's Entertainment Technology Center.

Panelists

Erick S. Dyke
n-Space

Richard D. Lindheim
USC Institute for Creative
Technologies

Michael Macedonia
US Army STRICOM

Randy Pausch
Carnegie Mellon University

Scott Trowbridge
Universal Studios Islands
of Adventure

Research and Education in the Arts and Technology (CREAT). He has spent more than 20 years as a creative principal in developing experiences and environments for entertainment, marketing, and education for clients such as Universal Studios, Nickelodeon Studios, Paramount Parks, Walt Disney World, Sanrio, and Dollywood. His education includes a master of fine arts in design for film, and a bachelor of fine arts in design for theater from New York University.

Erick S. Dyke

“Major and obvious changes need to occur in the infrastructure of the game market to allow for creative diversity and technical innovation, which in turn will push the true potential of the marketplace. The obstacles are more financial and operational than technical. Arcades and location-based entertainment present a great opportunity as well as the biggest disappointment. Much needs to be learned from consumers of the home video game market. The new model will not be from military or entertainment. We need to rethink the entire structure of how people enjoy interactive entertainment, then build the financial model to make it happen.”

Erick S. Dyke is president and founder of n-space, creators of interactive video game hits such as Duke Nukem: Time To Kill, RugRats 1 & 2, and Die Hard Trilogy 2. He started his career with General Electric and Martin Marietta working on several “killing simulators.” After extensive experience with corporate and military simulation contracts, he was inspired to start his own creative video game company. Partners include GTInteractive, 3DRealms, Nickelodeon, and Fox Interactive.

Richard D. Lindheim

“The collaboration between artist and scientist is not a new idea. Throughout history, artists and writers have served as the inspiration for scientists to turn their imaginative and sometimes fanciful ideas into reality. From the submarine to the cellular flip phone, we are indebted to the creativity of the artist and the brilliance and dedication of scientists. But this transference has never before been systematically approached. That is the role of the Institute for Creative Technologies (ICT).

“It will not be an easy task. In addition to difficulties in communication, where one word has very different meanings in the entertainment and scientific communities, the problem-solving approaches of the two worlds are quite different, as well as their criteria for success. Hollywood and the videogame world are masters of illusion; it is the perception of reality, not reality itself, that is important. Science must deal with reality itself and artistic “license” or deviation is not an acceptable alternative.

“This is only one of the challenges facing the ICT in its effort to create a more immersive, experiential, and educationally effective simulation training experience.”

Richard D. Lindheim is in charge of the strategic direction of the Institute for Creative Technologies (the US Army’s research program in entertainment simulation technologies, in conjunction with the University of Southern California’s Entertainment Technology Center, ISI labs, and the Annenberg Center). His work includes coordinating projects with the Hollywood-based entertainment industry, the computer game industry, computer technology companies, and related activities. His goals are to propose new projects that will integrate creative elements from the Hollywood and computer game industries with developing computer science technologies to provide more immersive and effective training applications for the military.

His affiliations include the Writers Guild of America, the Academy of Television Arts and Sciences, and the AFS Intercultural Programs, USA. His career began in the story department of CBS Television. More recently, he was executive vice president of the Paramount Television Group. He has also held executive creative positions with Universal Studios and the National Broadcasting Company.



Duke Nukem, DangerGirl, Rugrats, and DieHard Trilogy 2 from n-space interactive game developers.

Emotional Simulator: The Tears and Fears of Creating a Compelling Simulated Experience for Both Entertainment and Training

Michael Macedonia

“The computer and Internet revolutions have substantially changed the direction of entertainment from delivery via a mass medium such as television to a mass-customized experience via the Web and PC. However, the art of entertainment still requires stories, characters, and direction to make the experience meaningful and enjoyable. The US Army faces similar challenges as it adapts to changes brought about by mass marketing of supercomputing (for example, Playstation 2), low-cost graphics, and higher expectations of technically savvy soldiers.

“However, a more fundamental need is to simulate new types of problems such as urban conflict and peacekeeping. As witnessed on the bridge of Mitrovica, Kosovo, young men and women are often called upon to make critical, stressful decisions that have far-reaching political impact. Today, these scenarios are difficult to represent in virtual environments. The ability to employ story, emotion, and synthetic characters will enable us to represent the human dimensions of war and conflict to provide training for the truly difficult decision-making problems our soldiers must face.

“Establishment of the Institute for Creative Technologies is just one of many steps needed to provide the essence of verisimilitude in training and virtual reality systems. We will explore all avenues of entertainment technology to keep pace with the challenges presented to us, whether in application to distributed learning or embedded training systems. Ultimately, we want to prepare ourselves for the future by experiencing it.”

As chief scientist and technical director for STRICOM, Michael Macedonia is responsible for planning, development, coordination, and direction of a comprehensive program that provides the full spectrum of commercial and defense technology support for an \$800 million per year enterprise. He chairs the STRICOM Technical Council, which provides outreach to the commercial simulation industry.

A graduate of West Point, he served as an infantry officer in a number of United States and overseas assignments including Germany and the Middle East. He is a veteran of Desert Storm. Following his military service, he became vice-president of the Fraunhofer Center for Research in Computer Graphics, Inc. in Providence, Rhode Island. He then joined the Institute for Defense Analyses in Alexandria, Virginia.

Randy Pausch

“Using technologies like virtual reality to do ‘simulations’ is like using a movie camera to film a stage play. Sure, you can do it, but the really interesting stuff is in what you can’t do in a mindset of ‘simulating the real world.’ And, yes, you can bet that the entertainment industry will push on fantasy harder than, say, the military. The real problem with innovating toward the impossible is that it requires putting together artists and technologists: not an easy job.

“At Carnegie Mellon, our Entertainment Technology Center is focused on how to most effectively combine very different cultures: the arts and engineering. We are attempting, as we believe appropriate for a university, to address the fundamental academic questions that underlie this shift, such as: How can we effectively share control between an ahead-of-time author and a real-time audience, user, guest, or player? Ironically, those who have the most experience in the traditional media have both the biggest advantages and the biggest disadvantages when authoring in new media. They have great instincts from decades of experience, but those instincts do not always transfer well, especially from non-interactive to interactive media. At the Entertainment Technology Center, we focus on putting together students from widely varying disciplines and helping them learn the vocabulary, working styles, and values of other disciplines. In this collaborative environment, we can create new kinds of virtual environments, which we hope will be much more than ‘simulations.’ The real design challenge is building a world that ‘feels real’ but adds just enough magic so that it’s obviously not a simulation of any reality.”



The Spiderman Ride at SuperHero Island in Universal Studios Islands of Adventure.



Stan, the Human Patient Simulator, being treated at the Mass Casualty Training at UCF's Institute for Simulation and Training.

Randy Pausch is professor of computer science, human-computer interaction and design at Carnegie Mellon University, where he is co-director of the Entertainment Technology Center. 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, particularly for the DisneyQuest virtual-reality-based theme park. At Carnegie Mellon, he teaches the Building Virtual Worlds course, in which teams of artists, architects, designers, dramatists, computer scientists, and engineers work together to build full interactive, HMD-based VR experiences on a two-week, end-to end production schedule. He also leads the team that produces the Alice virtual reality authoring tool (www.alice.org).

Scott Trowbridge

“The landscape of entertainment has changed, and entertainment industry leaders need to not only recognize that change, but also embrace it to remain relevant to today’s and tomorrow’s audiences. With emerging technologies and Internet-based entertainment, an investment in developing new and compelling simulation technologies

is a key way to attract an audience with ever-widening choices. These new entertainment choices spring from a variety of new and non-traditional sources, and entertainment executives need to rethink the concepts of competition, R&D, and development time or they risk becoming empty brands compared to more innovative and aggressive companies.”

Scott Trowbridge recently completed one of the most advanced, state-of-the-art theme park rides using the most sophisticated simulation technology. He is presently the senior show producer for Universal Creative, which develops future attractions for Universal Studios. He is described in the Wall Street Journal and Wired as a “mastermind,” but, perhaps more appropriately, in The New York Times as “a bit of a geek.” He has spent most of his career dreaming up unique ways to entertain people. Ranging from film, theater, restaurant, and retail to state-of-the-art themed attractions, he has produced and integrated entertainment projects for clients all over the world. A graduate of the University of Southern California’s film school, he now looks for opportunities to combine aspects of various media and technologies to create new entertainment experiences.



Immersive virtual environments from USC Institute for Creative Technologies.

Movies, the theater, and various types of performances have contributed to our lives as both entertainment and art for a long time. Recently, there has been a trend to introduce interactivity into these media. One of the main forces leading this trend is the introduction of digital and computer graphics technologies that make it possible to create autonomous characters with either realistic or exquisite qualities. Another force is that interaction technologies such as speech recognition and image recognition have advanced so much that they can be applied to real-time human-computer interaction. Integration of these technologies has provided computer characters with interactive capabilities, and these are the key factors in the new interactive media.

On the other hand, this trend raises several questions. Why is it necessary to introduce interactivity into these media? Is “being” on the Titanic more interesting than “watching” the movie “Titanic?” Does interactivity really create something new and impressive in virtual environments? If so, what are the new experiences that can be created by interactivity? What types of techniques and technologies can be utilized or have been developed to realize this? Although these are crucial questions, there have been few discussions on these essential issues.

The goal of this panel is to examine these questions through the work of researchers who have been deeply involved in the actual making of new interactive environments and characters. Given the recent progress in technologies such as gesture recognition and speech recognition, we believe that now is a good time to discuss these essential topics with researchers, engineers working in these fields, and artists. We hope that several key issues for further research will be clarified through this panel.

Rodney Berry

“In my opinion, the purpose of interactivity is to increase the level of engagement in a virtual environment or multimedia work. After all, a microwave oven is ‘interactive,’ but it is not particularly engaging. A good painting, however, can be deeply engaging. In the latter case, interaction happens within the minds of viewers as they respond to the visual information.

Perhaps our idea of interactivity should extend beyond the sphere of the human-machine interface to include all forms of interaction: between humans and their own thoughts, between different humans, between parts of programs. Interactivity could be redefined as ‘capacity for interaction.’ The important thing to remember is that interactivity on its own does not guarantee engagement. Engagement comes from a feeling of personal involvement in the interaction. If I am asked to make choices, I need to feel that my choices will actually make a difference. Ultimately though, it comes down to how much our mind gets involved.”

Rodney Berry is currently a visiting artist/researcher at ATR’s Media Integration and Communications Laboratories in Kyoto, Japan. He holds a masters degree in media art from the University of New South Wales College of Fine Arts in Sydney, Australia. He has created many sculptural musical instruments and sound installations. He typically combines high and low-end technologies to integrate sound, visual, and spatial elements into a cohesive whole. The works often attempt to convey some sense of being in the presence of a living thing. More recently, he has focused on the use of computer-based artificial life to produce musical environments.

His work has been seen and heard in a number of Australian and international art events including the Third International Symposium on Electronic Art in Sydney, SoundCulture in Sydney and Tokyo, the Australian Perspecta exhibition in Sydney, the Third Australian Sculpture Triennial in Melbourne, Sculpture ‘85 in Melbourne, and Geographies of The Ear, a solo exhibition at The Performance Space in Sydney. His musical environment, Feeping Creatures, has been shown at Alife6 in Los Angeles, Virtual Worlds ‘98 in Paris, and the European Conference on Artificial Life in Lausanne, Switzerland.

Larry Friedlander

“All art is interactive in some sense. Digital media introduce new and transformative possibilities not anticipated by earlier media, yet it is foolhardy for exploration of interactive media to proceed without the informing presence of past examples, and an adequate theory of medium-user relationship and some typology of art forms. Once the novelty of interactive experiences wears off, users will demand from digital media a density and subtlety of meaning that we expect from all art. By a careful analysis of the connections and disjunction between old and new, we can devise strategies that add new dimensions to our experiments. These considerations can open up neglected avenues in design of interactive experiences and help move the field toward a sophisticated narrative and psychological level that it demands if it is to compete successfully with older forms.”

Larry Friedlander is a professor of English literature and theater at Stanford University, where he specializes in Shakespeare and performance. He is also co-director of the newly formed Stanford Learning Lab, a center for research and dissemination of learning technology and innovative curricula. He has been heavily involved in museum design and planning, and is now advising the Museum of Scotland, a new national museum in Edinburgh. He was an Osher Fellow at the San Francisco Exploratorium, where he helped design their new Learning Center. He has worked in major research laboratories such as the Apple Multimedia Lab and the Interactive Cinema Group at the MITMedia Lab, where he collaborated on a computer-enhanced theater space and narrative piece called the “Wheel of Life.”

Panelists

Rodney Berry
ATR Media Integration
and Communications
Research Laboratories

Larry Friedlander
Stanford University

Donald Marinelli
Carnegie Mellon University

Claudio Pinhanez
IBM T.J. Watson
Research Center

Donald Marinelli

“Audience empowerment means the ability of a user to become a genuine protagonist within a story. Truly interactive environments must transcend the current state of forced-choice and decision-tree paradigms that are at the heart of most interactive environments today. The computer-user interface must permit a lifelike interaction, wherein a user and computer graphics persona can converse in a non-linear manner, discuss the given circumstance, and establish and pursue objectives, while heightening the dramatic impact of the story.

“Carnegie Mellon’s Entertainment Technology Center, aware of the lifelike, invisible interface that interactive entertainment must achieve, has developed Synthetic Interviews based on the above considerations. Synthetic Interviews allow users to interact via spoken language with a computer graphics persona. The goal is an open-ended, non-linear dialogue, with the computer graphics persona deciding the parameters of what can be discussed. Since conversation is the basis of how most human interaction is conducted, Synthetic Interviews explore a primary way in which users become protagonists in an interactive environment.”

Donald Marinelli is co-director of the new Entertainment Technology Center, a joint endeavor between Carnegie Mellon University’s School of Computer Science and the College of Fine Arts. He is also a tenured professor of drama and arts management in the College of Fine Arts. Prior to accepting the position in the Entertainment Technology Center, for 14 years he was the associate head of CMU’s Drama Department, where he teaches theater history, critical writing, and arts management. He is also editor of the Dramaturgical Sourcebook Series developed by CMU Press and Smith & Kraus Publishers. He co-founded CMU’s master of arts management program, as well as the joint MFA graduate acting program between Carnegie Mellon University and the Moscow Art Theater School-Studio of Russia.

Ryohei Nakatsu

“I have been engaged in researching interaction technologies such as speech recognition, gesture recognition, and facial expression recognition. Also, I have studied the development of interactive virtual environments by integrating these interaction technologies with computer graphics technologies. This research confronts a big question: What will this integration create? Although many researchers and artists are working in this area, there is still no clear answer to this question.

“My answer is: ‘deep immersion.’ I believe that in the well-designed interactive virtual world, we would feel deeper immersion or involvement than we feel when we read novels or watch movies. Our basic attitude when reading novels or watching movies is passive. In the interactive environment, on the other hand, we can actively experience the virtual world by interacting with the entities of the world. This is the source for creating deep immersion or in other words active immersion.

“Of course, it is difficult to create this sensitivity by simply introducing interaction capabilities into the virtual environment. Therefore, the next question is: What is the key to realization of deep immersion? There are several key factors. In this panel, I clarify these key factors.”

Ryohei Nakatsu received his BS, MS, and PhD degrees in electronic engineering from Kyoto University in 1969, 1971, and 1982, respectively. After joining NTT in 1971, he mainly worked on speech recognition technology. Since 1994, he has been with ATR (Advanced Telecommunications Research Institute) and currently is the president of the ATR Media Integration & Communications Research Laboratories. His research interests include emotion extraction from speech and facial images, emotion recognition, nonverbal communications, and integration of multi-modalities in communications. He is a member of the IEEE, the Institute of Electronics, Information and Communication Engineers Japan (IEICE-J), as well as the Acoustical Society of Japan.

He organized courses for SIGGRAPH 96 and SIGGRAPH 97, and he organized and moderated panels for the 1998 ACM Workshop on Technologies for Interactive Movies and the 1998 IEEE Workshop on Computer Vision for Virtual Reality Based Human Communications.



From “It / I,” MITMedia Laboratory, 1997.
Photo by Claudio Pinhanez

Claudio Pinhanez

"In most video games and interactive installations, the action happens on what I believe is the wrong side of the screen: inside the computer graphics world. In my view, this is an unnecessary impoverishment of the experience, depriving users, and especially children, of opportunities for physical and social interaction with the real world and with real people. My work has been focused on how to create physical spaces where people can meet and play with computer graphics characters, objects, and landscapes in a real, physical space. Instead of virtual reality, I am trying to make fantasy more real.


"Through examples, I will show that some interaction technologies based on perceptual computing and tangible media are already providing the tools to build such interactive, immersive physical spaces. For instance, in 'The KidsRoom,' built in 1996, we used computer vision to track children and recognize dance steps in a complex narrative happening inside a real children's room. In 'It/I,' a computer theater play, I employed a stereo camera system to track an actor on stage and control his interaction with a computer graphics character.

"My research has convinced me that interactive stories, especially in physical spaces, should not be based on choice mechanisms but instead on exploration of local 'pockets' of interaction that do not disrupt the suspension of disbelief. Also, I contend that to achieve immersion in physical interactive spaces it is not necessary to resort to cumbersome devices like 3D goggles or highly artificial spaces like the CAVE. Instead, people can be entranced by carefully designed and produced stories and characters that, coupled with the right amount of unencumbered interactivity, respond to the user in the appropriate context."

Claudio Pinhanez is a media artist and computer scientist. Born in Brazil, he received his PhD in 1999 from the MITMedia Laboratory, where he conducted research on computer vision and artificial intelligence, and created and produced computerized performances. His work aims to create interactive spaces where users and performers can experience complex narratives while interacting with computerized actors. He coined the term "computer theater" in 1996 to describe the new and emerging experiences involving computers in theater, and since then he has been actively creating live performances inhabited by computer-controlled actors. He is also a member of the team that created the original prototype of "The KidsRoom," an interactive, immersive environment for children currently running in the Millennium Dome in England. He has also been a visiting researcher at ATR Laboratories (Kyoto, Japan) and the Sony Computer Science Laboratory (Tokyo, Japan), where he developed the "HyperMask" project presented in SIGGRAPH 99's Emerging Technologies. He is currently a research scientist at IBM T.J. Watson Research Center in New York.

Naoko Tosa

Naoko Tosa is a media artist and researcher in the ATR MIC Laboratories. She received PhD in art and technology research from the University of Tokyo, where she focused communication and used computers and electronics to create artwork that relates to the intelligence of emotions, consciousness, and unconsciousness. She specializes in creation of interactive art. Her best-known is was the Neuro-Baby project. Her work has been exhibited at the Museum of Modern Art, New York; the New York Metropolitan Art Museum; the annual SIGGRAPH conference; ARS Electronica; the Long Beach Museum; and other locations worldwide. Her works are also included in the collections of the American Film Association and 10 Japanese Art Museums.



From "Interactive Theater," ATR MIC Laboratories, 1998.
Photo by Naoko Tosa and Ryohei Nakatsu

Virtual reality (VR) technology has undergone a transition in the past few years that has taken it out of the realm of “expensive toy” and into that of functional technology. It integrates real-time computer graphics, body-tracking devices, visual displays, and sensory input devices to immerse a participant in computer-generated virtual environments (VE) that change in a natural way with head and body motion. Much as an aircraft simulator tests and trains piloting ability under a variety of controlled conditions, VEs can be developed to present simulations and augmentative stimulus inputs that are useful for cognitive and behavioral mental health applications.

In this panel, some leading VR and mental health scientists discuss their use of VR to assess and treat persons with phobias, post-traumatic stress disorder, burn- and cancer-related pain, traumatic brain injury, attention deficit and hyperactivity disorder, learning disabilities, and assessment of cognitive and functional abilities in aging populations.

Skip Rizzo

“My interest in virtual reality as a tool for mental health assessment and treatment stems from my previous work as a psychologist developing and implementing cognitive rehabilitation programs for clients with traumatic brain injuries. I saw first hand the limitations that exist in this area when using the traditional state-of-the-art tools. In fact, it is my view that much of the current assessment and treatment forms in psychology are still rooted in circa-1950s approaches. This is not so much an indictment of the mental health field as it is an observation of the limitations of the available technology in this area. Specifically, my work is guided by the view that well-suited matches exist for certain VR applications in the cognitive assessment and rehabilitation domain, as well as in the general mental health arena!


“Never before have we had the practical capacity to create mental health applications with the power to systematically deliver and control complex and graphically rich stimulus environments that also allow for precise measurement of behavior. These VR-mental health tools are a psychologist’s “dream” and, potentially, a patient’s deliverance! However, I believe that the key to success is in scientific determination of relative value of VR-based techniques compared to existing methods. To become enamored with these tools to the point of technological overkill could serve to drain resources from well-thought-out approaches based on good science. It is these views that drive my passion for developing and honestly evaluating new information technology tools that could make the best use of available resources for addressing mental health concerns.”

Albert “Skip” Rizzo has joint faculty appointments with the University of Southern California Integrated Media Systems Center and the USC School of Gerontology. He is the former program coordinator for the USC Alzheimer’s Disease Research Center and is the director of the Memory Enhancement Seminars for Seniors program. He is also director of the IMSC Virtual Environments Lab, which designs, develops, and evaluates the usefulness, feasibility, and efficacy of virtual reality systems that target assessment and rehabilitation of spatial abilities, attention, and other cognitive functions. He is also conducting research on computerized facial recognition, facial avatar animation, and better human-computer interaction systems for the elderly and persons with disabilities. His other IMSC activities involve provision of human-factors input on teleimmersion projects and research on integrating immersive audio in virtual environments.

He is the associate editor of the journal “CyberPsychology and Behavior” and is on the editorial boards of “The International Journal of Virtual Reality,” “Cognitive Technology,” and “Presence: Teleoperators and Virtual Environments.” He is the creator/manager of the Virtual Reality Mental Health Email Listserver. Prior to his USC affiliation, he was a cognitive rehabilitation specialist for eight years at the Coastline Traumatic Brain Injury Program in Costa Mesa, California, where he developed and implemented cognitive rehabilitation programs for clients with traumatic brain injuries. He received his PhD in clinical psychology from the State University of New York at Binghamton.

Larry Hodges

“Virtual reality is a field that has always been much better known for its potential than for its successful applications. One of the few areas where virtual reality has moved from laboratory prototypes to commercially available systems is psychological therapy. Most prominent among these applications has been the use of virtual environments in treatment of anxiety disorders. Hospitals and medical clinics in approximately a dozen locations in the US routinely use virtual reality exposure therapy to treat anxiety disorders such as fear of heights, fear of flying, and post-traumatic stress disorder. Other areas, such as the use of VR for pain management and rehabilitation, are developing rapidly and will soon be commonly available for general use. Current VR technology, with all its limitations, is still ‘good enough’ to support useful applications whose primary purpose is to give users the impression of being in a certain kind of space (such as is needed for exposure) or to direct the attention of users to their virtual circumstances (such as is needed for pain distraction or rehabilitation). In my opinion, the use of virtual reality in the general field of clinical psychology and rehabilitation will become the ‘killer app’ of VR.”



Virtual Classroom for Attention Deficit Assessment Developed at the USC Integrated Media Systems Center.

Larry F. Hodges is associate professor in the College of Computing and head of the Virtual Environments Group at the Georgia Institute of Technology. He is also co-founder of Virtually Better, Inc., a company in Atlanta that specializes in virtual reality exposure therapy of anxiety disorders. His research interests are in software and algorithm development, experimental quantification, and application development for virtual reality systems. He is on the editorial boards of two journals (“PRESENCE: Teleoperators and Virtual Environments” and “Cyber-Psychology and Behavior”) and is a member of the steering committee for the annual IEEE Virtual Reality Conference. He received his PhD from North Carolina State University in computer engineering, a MS in Computer science from NCSU, a MA in religion from Lancaster Theological Seminary, and a BA with a double major in mathematics and physics from Elon College.

Hunter Hoffman

“My colleagues and I at the University of Washington are exploring the value of virtual reality in clinical (spider phobia) and medical (adjunctive pain control) applications. The success rate for treating specific phobias with traditional in vivo exposure (showing them real live spiders) is high. So why do we need VR? Unfortunately, few phobics (less than 15 percent) ever seek treatment. Results of recent surveys suggest that the vast majority of phobics would be more willing to seek VR phobia treatment than traditional in vivo treatment. Furthermore, some types of phobias (fear of flying, for example) can't be easily treated in vivo (renting a real commercial jet for treatment, for example).

“A number of medical procedures are routinely conducted under medication conditions that leave the patient in severe-to-excruciating pain during the procedure. Immersive virtual reality is unusually attention grabbing, apparently because of the unparalleled illusion it gives computer users (patients) that they are inside the computer-generated environment. Burn patients undergoing wound care receive the same opioid medication they usually get, so in cases where VR doesn't work, patients are typically no worse off for trying VR.

“Our projects have been driven by psychological theory (models of human attention and pain) and treatment principles (well-established clinical treatments for phobias), with a good knowledge of the limitations of the technology. In addition to treating clinical populations and actual burn patients, our team uses analog laboratory experiments with undergraduates (for example, brief induced-pain experiments and analog phobia treatment studies with quasi-phobic students). In these

very efficient studies, data collection is completed in only a few weeks, whereas the same study with clinical patients (for example, burn patients and severe phobics) may take many months or years.

“We have used analog laboratory studies to identify the strengths and weaknesses of VR. These studies have already accelerated our understanding of how VR influences the processing of pain signals and whether it is effective for females, an understanding that will help us exploit the full potential of VR analgesia. Other analog laboratory studies with quasi-phobic undergraduates have also allowed us to discover that adding tactile cues to virtual spiders dramatically improves clinical outcome of VR exposure therapy. Watching severely burned kids experience intense pain during wound care creates a real urgency to quickly figure out how to optimize VR analgesia. Large (brief) declines in burn pain and suffering are common using VR analgesia, and there is huge room for improvement in the technique (for example, increasing treatment duration and frequency). Encouragingly, the price of VR equipment is plummeting, which will increase accessibility. Altogether, numerous factors bode well for the future of VR for treating phobias and pain (and many other maladies).”

Hunter Hoffman studied memory at Princeton University and investigated human memory and attention with eyewitness and false-memory expert Elizabeth Loftus. He is a project manager at the University of Washington Human Interface Technology Laboratory, an affiliate member of the psychology faculty, and he studies the use of immersive VR to reduce pain during patient care at Harborview Burn Center. He also conducts controlled laboratory experiments that explore the relationship between the illusion of “presence” in virtual reality and how analgesic effectiveness can be maximized. In other projects, he is exploring use of VR exposure therapy for treating spider phobia, the value of adding tactile cues (position-tracked toy spiders) to increase treatment effectiveness, and using real (mixed reality) chocolate bars to add taste to virtual objects. He has published over 20 peer-reviewed manuscripts, including six in the “Journal of Experimental Psychology.” He participated in a SIGGRAPH 98 panel on virtual healing, and his research has been covered by CNN, Scientific American Frontiers, and BBC and Discovery Channel documentaries. He received the Japanese Government Research Awards for Foreign Specialists (MITI) in Science and Technology when he was a visiting scholar at the University of Japan. He has also served as an ad hoc reviewer for JEP:LMC, Psychological Sciences, Presence: Tele-Operators and Virtual Environments (SIGGRAPH 99), the “Journal of Applied Psychology,” and “ACM Transactions on Human Computer Interaction.” He is also a scientific advisor to VR researchers in Spain and Japan.

Dorothy Strickland

“While virtual reality technology is admittedly more compute-intensive than many others, today’s powerful PCs, gaming machines, and graphics cards make real-world VR at home practical and affordable. Even so, other than games, few applications are offered to the general public. There are many reasons for this, from a perceived low demand for such products to complex, hard-to-use software and the low survival rates of companies that provide VR hardware. Questions such as liability and patent rights make the issue trickier.

“To mature into applications that offer serious benefits for the average person, VR must form a more standard, coherent definition of why it is worth the effort and offer clean, finished products that are easily accessible and usable by everyone. Although this is not a trivial transition for a technology that has primarily been a research toy, the benefits would touch everyone. Reproducing traumatic situations for treatment or learning dangerous skills in safe settings are just two of the applications where VR can offer value that is presently obtainable in no other way. Many small children die each year in house fires because they hide when the alarm goes off and cannot be found in time. Crossing busy streets is now a part of most children’s lives. I am presently attempting to build worlds to help with these two learning needs, as well as other safety issues that most children, not just those with special needs, encounter. In a world that is growing more complex daily, VR can make life a little easier to understand. My goal is to have those worlds combined with proven teaching methodologies and delivered in a convenient, low-cost way to everyone.”

Dorothy Strickland is president of Virtual Reality Aids, Inc., a company funded by a Phase II Small Business Innovative Research grant from the National Institutes of Health. The grant is designed to transfer VR technology to a practical system that will help parents and teachers use VR at home and schools for individuals with special learning needs. She is also an adjunct faculty member in the Computer Science Department at North Carolina State University. She has published over a dozen papers and book chapters, and she has been involved in many panels on using VR for special needs learning, including panel participation at SIGGRAPH 98 and SIGGRAPH 96, when she chaired a panel. She was an invited speaker at the 1998 IEEE Virtual Reality Annual International Symposium and a guest editor for a special “Communications of the

ACM” issue on the use of VR for health care in 1998. She received the “CyberEdge Journal” VR Application of the Year Award in 1995, when she first published her VR research. In addition, her research has been reported in the worldwide press, including an Associated Press feature article. She received her PhD in computer engineering from North Carolina State University and a masters in computer science from the University of North Carolina at Chapel Hill.

Benjamin Watson

“In addition to its application in therapy and rehabilitation, I believe that VR and related technology will prove very important for cognitive assessment and measurement, both for clinical use and for basic psychological research. I have consistently found that even inexpensive, PC-based VR simulators can differentiate between the driving ability of persons with and without head injury. In related work, my colleagues and I have found that naming time, a well-known cognitive measure, is sensitive to familiar graphical manipulations. This implies that many of the measures and manipulations that are possible (often uniquely) in VR are sensitive to cognitive ability.

“One of the ways in which cognitive assessment has often been used previously is to assess the mental ability of the elderly, in particular to detect the early onset of Alzheimers disease and related sorts of dementia. To examine the potential of VR for use in this sort of assessment, my colleagues and I have performed a series of studies that have generated extremely promising results. In my opinion, the future will see a burgeoning application of VR in these applications, resulting in earlier and more effective diagnoses and new and previously unachievable basic scientific results.”

Benjamin Watson is assistant professor in computing at the University of Alberta. His dissertation focused on user-performance effects of dynamic-level-of-detail management. His other research interests include object simplification, medical and spatial applications of virtual reality, and 3D user interfaces. He is on the conference committee for the IEEE VR 2000 conference, and the program committee for the Graphics Interface 2000 conference. He earned his doctoral and masters degrees at the Georgia Institute of Technology’s Graphics, Visualization and Usability Center and his bachelors degree at the University of California, Irvine.



Scene from “Virtual Vietnam” for post-traumatic stress disorder treatment (courtesy of Virtually Better, Inc.).



Scene from “Fear of Flying” scenario for treating flying phobia (courtesy of Virtually Better, Inc.).

The popular press calls it a “movement,” an “initiative,” even a “revolution”! In spite of the hype, “free” software is moving out of the lab and into mainstream computing. The revolution began with the Free Software Foundation’s GNU (Gnu is Not Unix) Project and has been fueled by Linux. The Internet has provided a “free” mechanism for distribution and collaboration. Virtual communities form around the software, and people willingly share their code and/or their experiences with it.

The panelists present two approaches to “free” software: open source and shareware. Open source advocates provide the source code to the “customer,” while shareware proponents provide low-cost or free personal licenses and may charge a larger fee for commercial licenses. Shareware seldom includes source code.

“Free” software is like water. Water is free, but sometimes the consumer pays for transportation (municipal water) or packaging (bottled water) or quality (filtration). Large and small companies are embracing “free” software as a cost-effective, high-quality alternative to closed systems.

Companies like Cygnus and Red Hat are betting their futures on mass acceptance of “open” software. These companies support, enhance, and repackage open source software. They have defined business models that include training, distribution, and ease-of-use enhancements. Other companies are converting products into open source (IBM DX, Netscape Communicator) or shareware (Qualcomm Eudora).

- What motivates people and companies to give up a source of income?
- Or is the software just not worth anything anyway? A loss leader?
- Are the developers on ego trips or really providing a service to computing?
- Does free software present a threat to companies that sell competing software?
- Or is the free stuff of lower “quality?”
- How free is free software? What are the hidden costs (training, maintenance, productivity, etc.)?

The SIGGRAPH community is a major consumer and producer of “free” software. In this panel, people who have developed graphics software for use by the SIGGRAPH community describe their philosophies, motivations, and business models. Their long-term goals range from having fun to making money to world domination.

Larry Gritz

“I am the author of the Blue Moon Rendering Tools (BMRT), which consist of a RenderMan-compliant ray tracer and related utilities. Unlike my fellow panelists, the source code for my software is not available at all, and the package is only free for academic and private use (but is exceptionally cheap even for commercial use). I’m often asked how I can charge so little, or given that I clearly don’t make money from it, why charge at all and why not distribute source?”

“Consider an artist who paints for pleasure but exhibits her work and affixes a \$100 price tag to a painting. Clearly, \$100 doesn’t cover the artist’s time at any reasonable rate, yet the painting is still not given away. Why not? The \$100 is a paltry sum, the sale is an exchange that creates a psychological “contract.” Somebody who pays even a token amount will feel obligated to take good care of it. I have a similar theory about software, so I charge a token amount to ensure that it finds a good home.

“Its distribution and price are only some of BMRT’s anomalies. BMRT is compatible with the ambitious RenderMan Interface Standard and Pixar’s flagship rendering product, PhotoRealistic RenderMan. As such, and since BMRT’s feature set is slightly different than PRMan’s (in particular, it can do ray tracing and global illumination), BMRT has become popular in production studios, and has been used on such films as ‘A Bug’s Life,’ ‘Stuart Little,’ and ‘Woman on Top.’ Why does my employer allow me to continue developing and distributing a piece of software that, at least superficially, appears to compete directly with its own? Or do they benefit somehow?”

Larry Gritz leads the “next-generation” rendering research team at Pixar and is reigning chief architect of the RenderMan Interface Standard. He has been at Pixar since 1995, developing new rendering techniques and systems, and periodically serving as a technical director. His film credits include “Toy Story,” “Geri’s Game,” “A Bug’s Life,” and “Toy Story 2.” His research interests include global illumination, shading languages and systems, and rendering of hideously complex scenes. Prior to joining Pixar, he authored the popular “Blue Moon Rendering Tools,” a shareware RenderMan-compliant renderer which supports ray tracing and radiosity. He has a BS from Cornell University and MS and PhD degrees from the George Washington University. He is co-author of “Advanced RenderMan: Creating CGI for Motion Pictures” (Morgan-Kaufmann, 1999).

Panelists

Larry Gritz
Pixar Animation Studios

Bill Hibbard
University of Wisconsin

Will Schroeder
Kitware, Inc.

Michael Tiemann
Red Hat, Inc.

Bill Hibbard

Respecting programmers and building software communities.

“The programming culture has problems, which are wonderfully illustrated in the Dilbert cartoons. Too often decisions are made by managers who do not understand software and do not respect programmers. Too often programmers hide behind obscure and secret knowledge and are reluctant to take responsibility for their own business affairs. The open source movement is primarily about creating healthy software communities, where programmers share knowledge, and assume responsibility and authority over their own work. Freely open source enhances the value of programmers, the people who understand software, and reduces the importance of people who control money.

“Creating open source software is lots of fun for programmers. You get pen pals all over the world, and some may pay your expenses to visit them. People lend helping hands, porting your software to all sorts of platforms, fixing bugs, and adding new features. Energetic people will produce whole new software systems based on your software. Well-connected people may send you photos of the Pope looking at a demo of your software. People will often offer to hire your services to customize your software for their needs, so there is a viable business model for open source software. Redhat, Precision Insight, and others are showing how to build whole businesses around open source.

“As the title of this panel implies, purchase price is only one component of software cost. The principal cost is developing local expertise with a system, and open source can give businesses confidence that this investment cannot be threatened when a system is no longer available. Furthermore, open source systems usually develop communities of experts who provide support. In my experience, newsgroups and mailing-list support are generally much better than 800-number support.”

Bill Hibbard is the leader of the SSEC Visualization Project at the University of Wisconsin - Madison, that created Vis5D, VisAD, and CavesD (Mesa was written by Brian Paul while he worked for the SSEC Visualization Project). These systems, including source code, are all freely available from www.ssec.wisc.edu/~billh/vis.html. Vis5D is widely used to visualize environmental simulations, and VisAD is a Java component library for interactive and collaborative visualization and analysis of numerical data. Hibbard has been a member of the program committee of the IEEE Visualization Conferences since their inception in 1990. He also writes the VisFiles column in Computer Graphics, the SIGGRAPH newsletter.

William Lorensen

William Lorensen is a Graphics Engineer in the Electronic Systems Laboratory at GE's Corporate Research and Development Center in Schenectady, New York. He has over 30 years of experience in computer graphics and software engineering. He is currently working on algorithms for 3D medical graphics and scientific visualization. He co-developed (with Harvey Cline) two popular isosurface extraction algorithms: marching cubes and dividing cubes. His other interests include computer animation, information display, and object-oriented software tools.

He is the author or co-author of over 60 technical articles on topics ranging from finite element pre/postprocessing, 3D medical imaging, computer animation and object-oriented design. He is a co-author of “Object-Oriented Modeling and Design” published by Prentice Hall, 1991. He is also co-author with Will Schroeder and Ken Martin of the book “The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics” published by Prentice Hall in November 1997. The text describes vtk, an open source visualization C++ class library. He gives frequent tutorials at the annual SIGGRAPH and IEEE Visualization conferences.

He holds 26 US patents on medical and visualization algorithms. In 1991, he was named a Coolidge Fellow, the highest scientific honor at GE Corporate R&D. Prior to joining GE in 1978, he was a mathematician at the US Army Benet Weapons Laboratory where he worked on computer graphics software for structural analysis. He has a BS in mathematics and an MS in computer science from Rensselaer Polytechnic Institute.

William J. Schroeder

“There is no such thing as free software. Even the simple acts of maintaining a Web site, downloading code, and storing it on disk cost something. Add to that the cost of developing, maintaining, and using software, and it's clear that significant costs accrue. So why, and how, is it that software can be given away free of charge?

“The answer is as complicated as human behavior. Business professionals see free software as a loss leader, giving away complex technology for which they expect compensation: training, support, consulting, maintenance, or to gain market share. Technologists and artists see it as a work of art, a changing expression of their understanding of technology and a demonstration of their talents. The altruistic give software away to serve their fellow humans. Extroverts do it to make friends and build communities. Freedom fighters do it to create ‘intellectual equity’ independent from any given company, employer, or organization. And leaders use these many motivations to work with others to build something they could never do by themselves.

“The final answer is most likely a combination of these and many other reasons, resulting in the variety of software distribution models now prevalent: open source, freeware, no-cost license, shareware, commercial, and so on. In addition, hybrid models are also important because, for example, they provide open source software in combination with proprietary products and services.

“Because of this wealth of possibilities, now is a wonderfully exciting time for technologists. We face the freedom of and responsibilities for creating experiences that are consistent with who we are and what we believe. And if we’re fortunate, we’re having fun and making a living along the way.”

William J. Schroeder is president and co-founder of Kitware, Inc., a small startup company that provides commercial support for the open source Visualization Toolkit software. VTK was initially developed as part of a text on visualization and is now used around the world in commercial, research, and teaching applications. In his current role at Kitware, he continues to develop the graphics, visualization, and imaging capabilities of VTK. Kitware is also building proprietary applications on top of VTK, such as the VolView volume rendering system.

He received his B.S. in mechanical engineering from the University of Maryland College Park in 1980 and obtained his PhD (while working full-time at the GE Corporate R&D Center) in applied mathematics from Rensselaer Polytechnic Institute in 1991. His current research interests include visualization, computational geometry, graphics, and numerical analysis. In 1992, he published one of the first papers on polygon decimation. His other areas of interest include swept-surface generation, stream polygons for vector field visualization, and object-oriented graphics and visualization architectures. He is a research faculty member at RPI and continues to explore the use of graphics in the computational sciences.

Michael Tiemann

Red Hat, Inc.

“The open source development model has proven time and again that it can produce software that is more reliable, resilient, and extensible than proprietary alternatives, but many people still believe that the model falls down in certain areas due to economic factors that still favor proprietary models. While not arguing that open source is a panacea (it’s not), I discuss where open source is winning (and losing) in the world of graphics today.”

Michael Tiemann is the CTO of Red Hat, Inc. Prior to joining Red Hat, he served as the co-founder and Acting CTO of Cygnus Solutions, which was acquired by Red Hat in January 2000. During his 10 years at Cygnus, he participated in a number of roles, ranging from president (from 1989-1995), director of business development (1995-1996), and director of technical marketing (1996-1998), to acting CTO (1999). Prior to co-founding Cygnus in 1989, he was a PhD candidate in the Stanford Graduate School for Electrical Engineering. Prior to Stanford, he worked at MCC (Microelectronic and Computer Technology, Corp.) in Austin, Texas, where he began making contributions to the open source software development community. His early work was focused on GNU software, which was in part an inspiration to Linus Torvalds back in 1991 and became the enabling technology for the creation of Linux. Today, he is a frequent speaker and panelist on open source software and open source business models.

It is a popular misconception that storytelling is exclusively a linear proposition. And while it is true that the most contemporary mass media (cinema, television, and the novel) are primarily linear, history presents us with a diverse array of interactive narrative genres. From mythological and religious rituals to the Greek Chorus, Commedia dell'Arte street theater, and the highly participatory theater in the round of Shakespeare's day, to narrative environments such as the Egyptian temples and medieval churches of pre-literate Europe, humans have been finding ways to interact with narrative for centuries. More contemporary forms include improvisational and street theater, interactive theater, and role-playing genres such as Dungeons and Dragons, the Renaissance Faire, and the Star Trek fan culture.

In exploring new directions for digital interactive storytelling, we ought to look "beyond the box." Well-worn genres with a proven track record of success, such as games, provide one direction. The emerging game-story hybrid, for example, has yielded a number of satisfying results, such as text-based MUDs and MOOs and highly successful role playing games, such as Ultima Online and EverQuest, as well as narrative adventure CD-ROM's and virtual reality games. Narrative environments like Myst have proved popular fare in the mass media in their use of a navigational, discovery-based narrative structure. The mystery or detective story-game hybrid has also proven highly successful in CD-ROM games such as Blade Runner and The X Files. Other directions can be derived from ingenuity and experimentation, combining engineering and art into new configurations. Artificial intelligence has emerged as a powerful next-generation storytelling tool for dynamic narrative experiences. And fine artists, who regularly invent their own genres, have found new and dynamic ways of integrating technology and the physical environment to create deeper and more rewarding forms of immersion in the narrative space. The notion of self-authoring, emerging from online communal storytelling environments, puts the author-as-God in a new and unique role of world-creator while the audience takes on the mantle of storyteller.

This panel convenes a diverse group of international experts from fine arts, entertainment, academic research, and business to discuss potential directions for the evolution of interactive storytelling. Each of the speakers combines an overview of theoretical ideas with examples of specific work that is setting new precedents for the interactive narrative genres of the future.

Steve DiPaola

Authoring the Self: Identity and Role-Playing in Virtual Communities

"As thousands upon thousands of people are daily inter-relating in avatar-based virtual communities, they are beginning to break down the old definitions of identity and self. Not only do they inhabit many identities that are each authentic in a particular context, but also they soon realize that identity itself is a dynamic construct that, like a leash, can be pulled in tight or given generous slack. In this way, one does not have to choose between the extremes of either playing a role or strictly being oneself, but instead can meander through the identity space of this role of the self.

"Our modern world has forced each of us to inhabit many roles in the course of daily life (with multiple business cards, emails, domain names, and titles that are associated with each of them). Virtual communities have taken this newly discovered need for multiple and differing identities and exploded the boundaries to the point where authoring the self is blurring the lines between non-fiction and fiction."

Steve DiPaola has been involved with computer-based characters for creating fiction and communities dating back to 1984, when he served as a senior member of the computer animation research group at the New York Institute of Technology. He specialized in 3D character animation R&D and animation production for film, TV, and his fine art work, specializing in 3D facial animation. He has also published several papers and book excerpts on the subject. From 1994, he became one of the original developers and creative director for OnLive, one of the first publicly available VR communities on the World Wide Web. He rejoined this project as director of development at Communities.com's OnLive Group, where he leads a team of artists, architects, user-interface designers, and musicians in designing and developing 3D avatars and virtual spaces. Prior to re-joining OnLive, he co-headed the San Francisco office of Saatchi and Saatchi's innovation arm, called Darwin Digital, as creative director. He has participated in Panels at SIGGRAPH 97 and SIGGRAPH 98, and in Technical Sketches at SIGGRAPH 99.



Self-authored characters convene and converse in OnLive's voice-driven online community.

Alex Mayhew

"Today computer games are at a similar stage to that of film in its early days. The games market is crammed with the equivalent of action movies, pandering to a young audience with unsophisticated and simplistic plots. Where are the games that can engage people emotionally beyond the realms of bloody carnage? After all, technology provides the potential for us to explore the length and breadth and true potential of our dreams, so why don't we do it?"

"The time is coming for a fresh perspective on computer gaming, one that allows for emotional engagement within a dramatic context. It is time for gaming for grownups. In my talk, I will present my thoughts on creating new genres of more sophisticated game-story-art hybrids that offer a higher level of both emotional and intellectual involvement. Because it is dynamic and participatory, interactive media should be able to provide even more satisfaction than traditional art forms, such as painting and literature. Since it incorporates aspects of both these disciplines, interactive media has the potential to engage and draw the user in at an even deeper and more meaningful level."

Alex Mayhew is a London-based new media artist who is best known for his work at England's renowned Real World Multimedia, including Peter Gabriel's EVE. More recently, he served as creative director and chief designer of "Ceremony of Innocence," an interactive CD-ROM based on Nick Bantock's Griffin and Sabine trilogy. It is a mysterious story of love at the edge of sanity, told through the intense and intimate correspondence of Griffin, a young artist from London, and his enigmatic South Sea Island muse, Sabine. "Ceremony of Innocence" was hailed as a watershed in the evolution of interactive art and narrative, and received 17 awards for artistic excellence and innovation, including Overall Winner of the Euro Grand Prix Multimedia 98, Best Of at the 1998 New Media Invision Awards in San Francisco, Best Product of the Year in Australia's 1997 ATOM Awards, and the International Award of Excellence in Canada's Atlantic Digital Media Awards in 1997.



The player helps the main character battle the monsters of his subconscious in Alex Mayhew's Dreamer.

He is presently working as an independent artist developing a new title, "Dreamer" (www.dreamer.org.uk), a game for grownups. "Dreamer" is a made-for-the-medium experience that takes place inside the head of a man called Max Reno, who is asleep in a hospital and will not wake up. The player is transported into Max's subconscious to discover the reasons for his condition and to attempt to wake him up. The whole piece functions on the basis of high action and emotional drama, in which the user plays a central role in creating the story experience.

Janet H. Murray

Dramatic Agency in Broadband Storytelling

"The long-anticipated convergence of television and computers is getting nearer and nearer. In light of the FCC mandate for digital broadcasting by 2006, television producers are displaying approach-avoidance behavior as they anxiously enter a domain in which the viewer will 'lean forward' and 'interrupt' their linear stories. This presentation offers a glimpse of the contested territories ahead and the kinds of television experiences that seem most promising for the new medium. It offers the concept of 'dramatic agency' as a key aesthetic principle of storytelling that exploits the new expressive potential of interactive TV."

Janet Murray began her career as a scholar of English narrative and published several volumes detailing the lives of Victorian women. This work grew out of her interest in what was left out of the great Victorian novels (what could not be said within the conventional narrative and cultural structures of that time and place). She was drawn to the new digital media because she saw its ability, like the book or the movie camera before it, to expand our ability to capture human experience, and because it holds the promise of expanding human understanding and human sympathies. In the late 1960s, she was trained as a systems programmer at IBM. In the early 1980s, bringing together her background in literature and computer science, she began working with advanced computing environments, inspired by the early work of MIT's Media Lab to initiate projects in humanities computing that formed some of the first explorations of educational "multimedia." In 1999, she moved to Georgia Tech, where she currently teaches in the Information Design and Technology graduate program, and is a member of the Graphics Visualization & Usability Center. Her main interest there is in interactive narrative, including digital television, virtual reality, videogames, and hypertexts, which she explores in her landmark book, "Hamlet on the Holodeck: The Future of Narrative in Cyberspace."

Panelists

Steve DiPaola
Muse Communications

Alex Mayhew

Janet H. Murray
Georgia Institute
of Technology

Sara Roberts
California Institute
of the Arts

Michael Thomsen
Interactive Institute

Celia Pearce

Celia Pearce is a multimedia artist and designer who currently serves as an adjunct professor of interactive media and researcher at the University of Southern California's School of Cinema-Television and the Annenberg Center for Communication, respectively. She has authored a number of papers on interactive design, as well as "The Interactive Book: A Guide to the Interactive Revolution" (Macmillan.) She spent 16 years in the theme park industry, where she worked as a writer, conceptual designer, creative director, and project manager, developing a variety of innovative interactive attractions and exhibitions. She has creative-directed several highly-acclaimed projects, including Iwerks Entertainment and Evans & Sutherland's "Virtual Adventures: The Loch Ness Expedition," a story-based 24-player virtual reality game that won several awards and was showcased as an interactive presentation in the SIGGRAPH 94 Electronic Theater. She also creative-directed bliss.com for Blitcom and Cosmo/SGL, the first streaming VRML character to run over the Internet, featured in the SIGGRAPH 97 VRML SIG. Fine art projects include Second Nature/The Virtual Gallery (in collaboration with the Space Dynamics Laboratory), which was shown at SIGGRAPH 95, and Body of Light (with musical director Bob Rice), an interactive performance that has been presented at Electronic Café International and the Banff Centre for the Arts. She has appeared in SIGGRAPH conference panels in 1994, 1996, 1997, and 1998, and her work was shown at SIGGRAPH 94, 95, and 97. She also served as Panels Chair for SIGGRAPH 98, and has spoken at numerous conferences and taught interactive design throughout the US, Europe, and Canada.

Sara Roberts

"People are attracted to interactivity not for its multiple-choice aspects but its ability to represent presence. We have a primal attraction to things that share our own time frame, something that was achieved through mechanical automaton, and now through the CPU. Generating and organizing material in real time gives us that same uncanny feeling. For some time, my work has been concerned with the characteristic back and forth of relationships. I have been writing and programming idiosyncratic emotional engines, which in interaction with each other, or with a visitor, generate narrative of a sort that is unpredictable in its particulars, but reliably invokes a certain relationship. The physical settings for these interactions are intrinsic to the relationship, just as the body and the room and the distances between us are intrinsic to our own experience of others. In an age where much of the digital and interactive world is divorced from the physical (both body and environment), I have been exploring a form of storytelling that is related to both space and time, that integrates the temporal and the dimensional. Since these things are fully integrated in 'real life,' it makes sense to create story spaces that reflect and explore the way we experience events in actual time and space."

Sara Roberts received an MFA from the California College of Arts and Crafts, Oakland in 1988. Her thesis project, the interactive installation "Early Programming," was shown at the San Francisco Arts Commission Gallery and the Long Beach Museum of Art. She collaborated with artist Lynn Hershman on two interactive installations, Deep Contact (1990) and A Room of One's Own (1992), that were shown internationally. In 1993, she was a resident at the Djerassi Foundation, and her interactive sculpture Digital Museum was commissioned by SGL for the Interactive



Visitors navigate different viewpoints in the same scenario in Sara Roberts' installation *Elective Affinities*.

Gallery, Moscone Center, San Francisco. In 1994, she received a WESTAF grant for New Genres and completed *Elective Affinities*, a large-scale installation shown first at the Yerba Buena Center for the Arts in San Francisco. *Elective Affinities* was shown in 1998 in Helsinki, and in Clermont-Ferrand, France in 1999. In 1997-98, a solo survey of her work titled *Physical Fictions* went to Filmhus Denmark, in Copenhagen; the Arizona State University Art Center, Tempe; and the Williamson Gallery at the Art Center College of Design, Pasadena. A collaboration with the animator Christine Panushka, a friend since childhood, produced a double self-portrait called *Archive: Comings and Goings 1-49* for the *Interactive Frictions* exhibition at the University of Southern California Fischer Gallery in June 1999. Also *Silence*, a site-specific work in a limestone mine in Telluride, Colorado, was completed and shown. She has taught previously at San Francisco State University and The San Francisco Art Institute. She joined the faculty of the School of Music at Cal Arts in 1994, in the Composition/New Media Program. Since 1996, she has been director of the Integrated Media Program at Cal Arts, an institute-wide graduate program for interdisciplinary work with technology.

Michael Thomsen

Creating New Narrative Forms: Tradition and Transcendence in Narrative Theory and Practice.

“Over the last 15 years, the computer has developed into a powerful creative tool that involves and reshapes just about every creative discipline. As a medium, the computer promises to restore interaction to the arts, and as a tool it allows creative people to work in completely new ways that require both new skills and new disciplines. With the advent of interactive media and collaborative virtual spaces, we are forced to reevaluate the concepts of artist, producer, and consumer. We need to develop new theory as well as practice, and we must rethink traditional disciplines and their relation to each other.

“In the Narrativity and Communication Studio in Malmö, Sweden we combine tradition with transcendence and artists with technologists to explore new avenues in interactive media, breaking the traditional boundaries of academia, art, and craft. My presentation discusses the theoretical basis of our work there and gives some examples of how we are putting these ideas into practice in approaching the problem of creating new genres of narrative for the interactive age.”

Michael Thomsen is research director for the Interactive Institute, a national digital media research institution in Sweden, and former director of the institute's Narrativity and Communication Studio in Malmö. He earned an MSc in computer science and mathematics from Aarhus Universitet and has worked as a systems developer for many years.

From 1991 to 1999, he worked for the LEGO Group (from 1993 to 1996, as manager of the group's Boston-based technology R&D department, where he acted as a liaison with the MITMedia Lab, and from 1996 to April 1999, as director of software R&D). He is also Chairman of the Board for Space Invaders, an innovative multimedia school in Copenhagen, and member of the board of Nordic Interactive. His main interests are in creativity, digital structures, organizational development, and the blurring boundary between the virtual and the real.

From the earliest days of computer graphics, researchers and artists have been fascinated by the challenge of digitally reproducing the human face. More recently, they have speculated about the feasibility of replacing actors with digital characters.

Early attempts such as Philippe Bergeron's Tony de Peltrie and Kleiser/Walczak's Nestor Sextone set the stage for the "Experience Music Project." In this location-based entertainment experience, a group of artists and technicians from Digital Domain combined Paul Debevec's lighting research, science, and art to recreate James Brown as a dancing, singing, funk phenomenon. With a unique application of subdivision surfaces, inverse subdivision surfaces, new tracking techniques, and traditional animation, they created a digital prosthetic. This panel looks behind the process of integrating advanced research, practical production, and traditional art and animation to set new technical standards and provide a unique, incredible entertainment experience.

Mark David Brown

"The animation team for the Young James Brown sequence was challenged with designing a robust system that would bring traditional key frame animators and motion capture together in a flexible and meaningful way. From my past experience with motion capture on 'Titanic,' I knew you could not take the animator out of the process. In fact, the animator must remain an integral part of the pipeline. A system that does not include the animator as an artist is, in my opinion, flawed and reduces the experience to an inflexible technical execution that would never withstand the constantly changing environment of film production, nor would it render a flawless performance. In this endeavor, I feel we have designed a truly innovative system that uses the best parts of motion capture, while returning the control to the key frame animators. This control enhanced the performance and, in some cases, made up completely new dialogue and expressions through traditional key frame methods, resulting in a true marriage of art and science."

Educated in traditional art at the Art Institute of Chicago and the California College of Arts and Crafts, Mark David Brown has animated on a computer since he bought his first Amiga in 1992. Hired first as an animator at Pacific Data Images in 1993, he moved to Cinesite in 1995, where he worked on such films as "Free Willy II" and "Waterworld." From there, he moved on to the world of post production before arriving at Digital Domain in 1996. Since then, he has contributed to "Terminator 2-3D" and "Bud Ants 2." He supervised a General Cinemas trailer and various other commercials before beginning character research and development for "Titanic." As one of the first three artists assigned to the show, which was code named SOS, he and his team were challenged with the task of creating more than 1,000 computer-generated passengers on the ship. After accomplishing that feat, he was named digital stunt supervisor for a major sequence of shots. In recent years, he has been involved in establishing the Digital Domain Character Department, while continuing to work on films and commercials, including supervision on a music video featuring Busta Rhymes and Janet Jackson. He is currently the animation supervisor for the Experience Music Project, for which he is producing the first photo-real digital human face.

Andre Bustanoby

"A hybrid of a ride film and a music video, the Experience Music Project is the culmination of an artistic and technical challenge. As the visual effects supervisor, it was my job to help organize, plan, and execute, in collaboration with a world-class team of artists, a varied and sensible approach to the work's visual challenges.

"Under Ray Giarratana's direction, and in concert with Vince Toto's cinematography and Bradley Thordarson's production design, the visual effects transported the audience to new places and times, to the beat of the music that was fundamental to the core of our story. Given Frank Gehry's design of the theater space and working at 70mm (five-perf) with a screen size and curvature unique to this presentation, we have created some novel solutions, both artistic and technical.



A still frame from Paul Debevec's "Rendering with Natural Light" from the SIGGRAPH 98 Electronic Theater. The piece was the first to use image-based lighting; the illumination was captured in the University of California, Berkeley eucalyptus grove using high-dynamic photography and simulated using the RADIANCE synthetic imaging system. See also: www.cs.berkeley.edu/~debevec/RNL/

“At the cornerstone of the effects work was the requirement to generate a young James Brown, circa 1969-1970. Full realization of this character required state-of-the-art CG facial animation and 3D facial tracking of our dance double for Mr. Brown. We were able to preserve James Brown’s actual facial performance while allowing for critical animator contribution. Finally, we employed new lighting and rendering techniques to match our computer-generated young James Brown with the dance double, resulting in a seamless blend of real and virtual.

“Furthermore, the director required use of the ‘Frozen Moment’ technique to integrate with the flowing nature of the cinematography. Instead of circling around a central subject, we designed the shots to flow more like a camera on a crane or a Steadicam. This helped the shots feel more natural and less a separate, obvious visual effect.

“Also critical to the narrative was creation of a completely computer-generated journey emphasizing the essence of funk. Using new software tools to tap directly into the music, we leveraged the beat to drive the animation of literal and figurative aspects of our funk universe. We also created a completely photoreal digital backlot version of Paramount Studios’ New York street, from the point of view of a sweeping CableCam, to give a sense of place and time to the intrepid protagonists who take the journey into the origins and nature of funk.

“As always, our job was to serve the story and help the director realize his vision. Through collaboration and teamwork, this task has proven both challenging and artistically satisfying to all involved.”

Andre Bustanoby is a visual effects supervisor and the performance capture supervisor at Digital Domain. He received his bachelor of science degree in mechanical engineering from Washington State University in 1990. Upon graduation, he worked with Masters FX on “The Howling VI,” “The Resurrected,” and “Tales from the Crypt” for HBO. Later that year, he moved to Boss Film Studios, where he served as a design engineer. Over the next five years, he moved up to chief engineer and finally to head of integration. In that position, he supervised CG/live-action integration issues like 3D tracking, set and lighting survey, and motion-control integration with CG.

Late in 1995, he joined Digital Domain, where he was tasked with supervising and managing the then-nascent Performance Capture Department. After providing motion capture for such award-winning commercials as Nike’s “Virtual Andre,” Coke’s “Jitterbug,” and the Academy Award-winning feature film “Titanic,” he co-supervised the visual effects for Bjork’s “Hunter” music video. Later, he acted as the visual effects supervisor for the 1999 summer film “Lake Placid.” Currently, he is supervising the visual effects for Paul Allen’s “Experience Music Project.”

Paul Debevec

“Realistically combining live action and computer-generated imagery is one of the most important yet one of the most challenging feats required in today’s visual effects work. Perhaps the hardest part is matching the lighting between the CG objects and the real environments in which they are placed. This process has traditionally involved adjusting color balance, shading, highlights, and shadows by eye and by hand, often unsuccessfully and always with a great deal of effort. At the University of California, Berkeley, we’ve developed high-dynamic-range photographic techniques to record incident illumination in real environments and used global illumination algorithms to illuminate CG objects with these measurements of real light. The result is that CG objects can be realistically illuminated with precisely the light they would have received had they actually been present in the scene.

“We have used this image-based lighting technique to create a variety of animations, including ‘Rendering with Natural Light’ (SIGGRAPH 98) and ‘Fiat Lux’ (SIGGRAPH 99). However, it remains to be seen whether these techniques are adaptable and useful within traditional production pipelines. For the ‘Experience Music Project,’ Digital Domain made the leap of implementing these lighting techniques, as well several other cutting-edge techniques for modeling and rendering, with visually amazing results. The greatest result, I hope, will be a greater predisposition in the visual effects industry to experiment with and implement new techniques and in so doing to further advance the state of the art.”



*Mark David Brown
Animation Supervisor,
Digital Domain*



*Andre Bustanoby
Visual Effects Supervisor,
Digital Domain*

Panelists

Mark David Brown
Digital Domain

Andre Bustanoby
Digital Domain

Paul Debevec
USC Institute for Creative
Technologies

Joshua Kolden
Digital Domain

Paul Debevec is a research scientist and filmmaker at the University of California, Berkeley. His research combines digital photography with computer graphics, creating 3D camera moves from 2D photographs, illuminating synthetic objects with real light, and realistically combining real and computer-generated imagery. Using these techniques, he has directed several internationally exhibited computer animations, including "The Campanile Movie," "Rendering with Natural Light," and "Fiat Lux." which premiered in the annual SIGGRAPH Electronic Theater. Recently, his virtual cinematography techniques were employed at MANEX Entertainment in creating the "bullet time" shots for "The Matrix." His current work involves applying image-based lighting techniques to digital filmmaking, virtual actors, and sets, and documentation of archaeological sites. He received his PhD in computer science from the University of California, Berkeley in 1996.

Joshua Kolden

"Balancing exciting technical innovation with production-minded artistry is very difficult but vital. Technical innovations provide the foundation to take ideas further, and artistic innovation raises the bar for new research. As exciting as one aspect might be by itself, they are infinitely more successful when they are combined. What made this project possible and successful was the single-minded team effort to achieve one goal: Create a digital version of James Brown that is indistinguishable from Mr. Brown himself on camera.

"We did not allow ourselves the luxury of using our favorite software, algorithms, or animation techniques. We threw our previous experience aside and started from scratch, with a clean slate, nothing but the latest tools, and the best animators to choose from. Being

production minded, we were careful not to over engineer. Being quality minded, we were careful to include the best technology. My strongest hope for the people doing work in this field is to improve human communication. To this end, it is my experience that innovation will be achieved only by artists willing to embrace the new and researchers willing to look past the numbers. I am very proud to work alongside such innovative people."

Joshua Kolden's work as a semi-professional actor during high school served as a foundation for his current work in character development. In film school during the 1980s, he received his first glimpse of computer graphics on the Pixar II image computer. Armed with an advance copy of the RenderMan Companion, he started using RenderMan toward the end of 1988. In 1993, after working in broadcast for a few years, he was invited to teach a masters degree computer graphics course at the University of Arkansas at Little Rock and went on to help design the computer graphics curriculum there. He also consulted for a broad interdisciplinary project that provided high-level computer graphics equipment to other departments, including the art department, the film department, and computer science. Since 1995, he has worked on traditional visual effects. His credits include: "Free Willy II," "Waterworld," "Terminator 2-3D," "The Fifth Element," "Titanic," "Flubber," "Armageddon," "Stuart Little," and "Deep Blue Sea." He is currently directing a feature film that he wrote, which includes absolutely no computer-generated effects.



Paul Debevec
Research Scientist, University of
California, Berkeley



Joshua Kolden
CG Supervisor, Digital Domain

Technically and aesthetically innovative artists discuss their work, with no art jargon! Each speaker has work featured in the SIGGRAPH 2000 Art Gallery. They explore the nature of interactivity, art and science, cause and effect, installation and space, and what artists can learn from game designers. Their projects push forward the boundaries of technologies ranging from autonomous agents and artificial life to Web visualization and real-time interaction with 3D worlds.

THE ART AND SCIENCE OF BIOTICA

Richard Brown

"Biotica is both research project and art installation, carried out in the 'art as a mode of enquiry' territory in the Computer Related Design Research Department at the Royal College of Art. The Biotica project began in May 1998 as the result of two successful award bids: A Wellcome Trust Sci-Art award, in collaboration with Igor Aleksander, a neural net specialist from Imperial College, London; and a three-year award from Intel Corporation for development of interactive real-time 3D artworks. Biotica amalgamates the science of artificial life, 3D graphics, and immersion and interaction technologies with the art of abstraction and engagement, questioning ideas about perception, behaviour, emergence, and aliveness. The presentation uncovers the research history, technologies, prototypes and ideas, inspirations, and influences that underpin the Biotica installation in the SIGGRAPH 2000 Art Gallery."

For the last four years, Richard Brown has practiced art as a mode of enquiry to research creation of "virtual unrealities" within the Computer Related Design Department of the Royal College of Art, where he is a research fellow. This has led to creation and international exhibition of two interactive installations: Alembic and Biotica. His work is inspired by a mix of art, science, and technology. This hybrid approach is also reflected in his educational background: a BSc in computers and cybernetics and an MA in fine art. Prior to his research work, he worked in commerce, running his own company, Interactive Environments, which specialised in creative application of technology in industrial communications and entertainment.

DEEPER CONVERSATIONS WITH INTERACTIVE ART

Andrew Stern

"The past few years have seen a flurry of activity in computer-based interactive art, spurred on by the advent of the Internet and the World Wide Web. Although many artists are being very creative and innovative, I find myself searching for artworks that truly take advantage of the capabilities that, from my experience as a designer and programmer of interactive virtual characters, I know the computer can offer.

"My panel presentation takes a step back and looks analytically at the computer as a medium for art, in an attempt to understand the important new capabilities it offers to artists. First, I present the point that 'push-button' interactive art offers only a superficial form of interactivity, that artists should strive to deepen the "conversation" between art and audience. This discussion proposes some technical methods to achieve this deeper conversation. Secondly, I argue that the best game designers are addressing issues of concern to artists, and that by studying today's most compelling games, additional insight can be gained for creating more conversational interactive artwork. I conclude with the assertion that the creation of deeply interactive artwork fundamentally requires programming, and therefore most artists will need to collaborate with programmers, which raises important issues of authorship and control."

Andrew Stern is a designer and programmer of the interactive characters Babyz, Dogz, and Catz from PF.Magic in San Francisco. Along with his fellow creators Adam Frank, Ben Resner, and Rob Fulop, he has presented these projects at a variety of conferences, including Digital Arts and Culture 99, AAAI Narrative Intelligence Symposium 99, Autonomous Agents 98, and Intelligent User Interfaces 98. Catz won a Design Distinction in the first annual I.D. Magazine Interactive Media Review, and along with Dogz and Babyz is currently part of the American Museum of Moving Image's Computer Space exhibit in New York. The projects have been described in publications such as the New York Times, Time Magazine, Wired, and AI Magazine. Stern is currently collaborating with Michael Mateas on an interactive drama project. He holds a BS in computer engineering from Carnegie Mellon University and a masters degree in computer science from the University of Southern California.

HAME

Laura Beloff

HAME, an installation in the SIGGRAPH 2000 Art Gallery, investigates hysteria and boredom. The installation consists of two sculptural objects in space, stereographic video projected on the images, interactive jackets, and audio. This presentation discusses the work's conception, execution, and operation, as well as related work by the artist.

After completing Fulbright-supported studies at the California Institute of the Arts in 1998, Laura Beloff was a 1999 resident at the Ars Electronica Center.

STARRYNIGHT

Mark Tribe and Alex Galloway with Martin Wattenberg

“Each time someone reads a text at Rhizome.org, a dim star appears on a black Web page. When a text is read again, the corresponding star gets a bit brighter. Over time, the Web page comes to resemble a starry night sky, with bright stars corresponding to the most popular texts, and dim stars corresponding to less-popular ones. Clicking on a star triggers a special pop-up menu. The first menu option allows you to read the text that corresponds to the star. The second menu option allows you to select a keyword associated with that text. When a keyword is selected, STARRYNIGHT draws a unique constellation of stars whose texts share that keyword. You can use these constellations to find other related texts, and in doing so, follow your interests through the vast array of ideas and information in the Rhizome.org archive.”

Alex Galloway is editor of Rhizome.org. He has written on theoretical issues surrounding digital technologies and is currently developing new interfaces for navigating information on Rhizome.org. He studied semiotics and poststructuralism at Brown University and received a BA in modern culture and media. He is currently a graduate student in the literature program at Duke University.

Mark Tribe founded Rhizome.org, an online community for digital artists, and StockObjects (www.stockobjects.com), a new-media stock library. Prior to Rhizome and StockObjects, Mark worked as an artist in Berlin, and developed commercial Web sites at Pixelpark GmbH, a leading German interactive media firm. His art practice has included happenings, installations, and net-based work. Mark speaks widely on new-media art. Recent engagements have included appearances at Ars Electronica, the annual SIGGRAPH conference, ISEA, the Dutch Electronic Arts Festival, the European Media Arts Festival, Harvard University, Pratt Institute, and Brown University. Mark received a masters of fine arts in visual art from the University of California, San Diego (1994) and a BA in visual art from Brown University (1990).

Martin Wattenberg is a computer programmer who lives in New York.



Virtual Babyz (1999) from PF. Magic/Mindscape.

Noah Wardrip-Fruin

Noah Wardrip-Fruin is an artist, writer, and research scientist at the New York University Media Research Lab. His current projects include The Impermanence Agent, a storytelling Web agent that customizes its story of impermanence for each user. The customization continues until none of the original story remains. The Impermanence Agent is featured in the SIGGRAPH 2000 Art Gallery. His previous work includes “Gray Matters,” a collaborative fiction embedded within images of the human body, which he completed as an Edward Albee Foundation Fellow. Its exhibition at the Sandra Gering Gallery was the first public installation of a zooming user interface.

Diane Gromala

Diane Gromala is associate professor in the School of Literature, Communication, and Culture at Georgia Tech, where she teaches in the graduate program in Information Design and Technology. She is also a faculty member of the Graphics Visualization Center. Her courses focus on the cultural and theoretical aspects of new and virtual media. Gromala’s critical analyses of virtual media are informed by her work as an artist and designer in virtual reality and other forms of digital media. Her articles have appeared in numerous art and design journals, and her virtual reality artwork has been performed and presented in Canada, the US, Europe, and the Middle East, and aired on the Discovery Channel and the BBC. She is on the Editorial Board of Postmodern Culture, Chair of the SIGGRAPH 2000 Art Gallery, and a recent recipient of a Senior Fulbright Fellowship.



Screen shot of Starrynight (www.rhizome.org).

The growth of the Internet has been fueled by millions of Web pages, which primarily contain HTML-based text and simple graphic images. As the capabilities of web browsers are enhanced, more and more information is being loaded onto these pages in the form of rich media. Initially the rich media has tended to be enhanced graphics or animations (for example, animated GIF files), but that is changing with new rich media formats emerging (Shockwave, etc.) along with new techniques for using them. But is all this rich media a benefit or annoyance to the user browsing the Web?

This panel discusses the technology of rich media and how it can be expected to change the Web (for better or worse) in the coming months and years. It provides an overview of rich media technologies and demonstrates the use of graphics, sound, and animation to enhance Web sites, implement engaging games and activities to replace simple banner ads, impose interstitial/superstitial ads and transition screens on Web pages, and provide promotional materials that are entertaining and informative. The panel also explores the bandwidth consumed to provide these enhanced Web experiences and debates and reviews the pros and cons of various approaches.

Web technologies reviewed include:

Shockwave
Fireworks
Enliven
Onflow
MPEG-4
SMIL

Each technology is highlighted with a slide detailing its capabilities, characteristics, and market penetration (number of Web pages with the embedded technology). After the technology review, the panelists present their particular points of view regarding rich media.

RELATED INFORMATION

Wired Digital Rich Media Study
hotwired.lycos.com/home/advertising/rm/index.html

Interactive Week: Rich Media Ads The Talk Of The Town
www.zdnet.com/intweek/stories/news/0,4164,2155594,00.html

@Home Network Concludes Rich Media Ad Study
www.internetnews.com/IAR/article/0,1087,12_74731,00.html

Rich Media SIG
www.richmediasig.com/westcoastlaunch.html

Mickey W. Mantle

A member of the "University of Utah gang," Mickey W. Mantle was a contributor to many early computer graphics products from Evans & Sutherland and Pixar Animation Studios, including Pixar's RenderMan software. He joined Broderbund Software in 1991 as vice president of engineering/chief technical officer, where he managed technology initiatives and contributed to development of many successful and award-winning products, including Living Books, Myst, Riven, and many more.

In 1999, he joined CDDB as vice president of development where he oversees all development of CDDB Internet services. This includes software for enabling CD audio player applications such as MusicMatch Jukebox, RealJukeBox, and many others. He has a degree in computer science from the University of Utah and has been attending the annual SIGGRAPH conference regularly since 1978. He last chaired a SIGGRAPH panel in 1991.

Michelle Y. Kim

Michelle Y. Kim is a member of the MPEG-4 committee, editor of MPEG-4 Amendment 2, editor and AHG chair of MPEG-4 Textual Format, and a member of the W3C SMIL Working Group. At IBM T. J. Watson Research Center, she manages the Composite Media Technologies Group, whose focus in recent years has been to contribute to and shape emerging multimedia integration technologies for both Internet and broadcasting applications.

Panelists

Michelle Y. Kim
IBM T.J. Watson Research
Center

Diane Rogers
Macromedia, Inc.

H.B. Siegel
Amazon.com

Douglas Tuttle
Onflow Corporation

Diane Rogers

“Rich media infinitely expands the possibilities of the web. Shopping experiences become more like the real world, games become more engaging, presentations become more meaningful. As bandwidth gets bigger and bigger, and the technology becomes faster and faster, consumers can finally match the visceral values of television with the interactivity of the Web. Rich media is the enabler for a new era of communication.”

Diane Rogers is vice president of product management at Macromedia, Inc. In this role, she works with a team of engineers and product managers to design innovative software for Web publishing. Working with Flash, Fireworks, Director, FreeHand, and Dreamweaver, she participates in changing the face of the Internet. Prior to Macromedia, she led marketing and product management for Borland International’s programming environments Delphi and C++. For 10 years, she has lectured worldwide on emerging software trends. She has written training videos, authored technical articles for the press, and participates in technology conferences.

H.B. Siegel

“At Amazon, the challenge is to identify and deploy media technologies that truly add value to the customer experience. The Web customer base is extremely diverse, from very-low-bandwidth handheld devices to text browsers (they still exist!), to modems, and finally to broadband home and workplace connections. I will review some of the decisions that help guide large e-commerce companies in evaluating and integrating these technologies: examining the customer experience from many vantage points including the connection speed, first-mover advantage, marketplace perceptions, platform, portability, security, reliability and, of course, how enjoyable it is use.

“In addition, it is often challenging to find the right balance between partnering and acquisition, and where it may be advantageous to develop the technology internally.”

H.B. Siegel is in charge of Amazon.com’s newly formed Media Technologies group. Before joining Amazon, he served as chief technology officer for ILM, where his teams produced the technology enabling the production of dozens of films, including “Star Wars I” and “Jurassic Park: The Lost World.” As director of rendering for SGI from 1994 to 1996, he led teams that built SGI’s graphics application programming interfaces, including the industry leading OpenGL and Performer.

Previously, he led the development teams for several software initiatives including the influential Dynamation and Kinemation at Wavefront Technologies in Santa Barbara as director of animation and rendering. From 1984 to 1991, he had programming and management roles with Pixar Animation Studios encompassing development libraries for the Pixar Image Computer and interactive rendering. He has degrees in electrical engineering and computer science from Brown University and a masters in computer science from the University of California, Berkeley.

Douglas Tuttle

“Rich media can be a curse in low-bandwidth situations. File sizes are too big, and the solutions are too clumsy. New integrated media types built for a real narrow-band world will leverage new levels of CPU performance and improved compression technology to bring high-quality rich media experiences to everyone, everywhere. Advertisers will embrace a technology that lets them preserve their TV brands on the Web. This will pay for the content people really want.”

Douglas Tuttle has worked in computer graphics, video and film production, or graphics tool development for nearly 30 years. He is currently vice president of engineering at Onflow Corporation. His expertise includes 2D and 3D rendering algorithms, light transport physics, compiler design, hardware design, optics, and system architectures. He has worked for Industrial Light + Magic, Macromedia, Sun Microsystems, Synthetic Video, Digital Research, and Lawrence Berkeley Laboratory. He has combined his deep technical knowledge with his artistic background in painting, photography, and even 1960s light shows to design superior production tools for real-world artists. He has a degree in physics from the University of California, Berkeley.

We will never know precisely what was in the minds of our ancestors as they painted the rock shelters in France 25,000 years ago, or raised the pyramids in Egypt, or even purchased a particular brightly coloured pot during the Middle Ages. Instead we rely on archaeologists to help us interpret the material preserved from ancient cultures in order to have some understanding of the past.

In the past few years, archaeologists have been increasingly turning to computer graphics and interactive techniques to help with their interpretation. This had led to a growing number of collaborative projects between archaeologists and computer graphics specialists, considering diverse topics from creating detailed computer models of art treasures such as Michaelangelo's Pieta to reconstructing ancient environments including the pyramids, Stonehenge, and Mayan cities, and even to investigating the perception of prehistoric rock art under prehistoric illumination conditions.

Many of the techniques used are extremely expensive, well beyond the traditional budgets of the archaeologists. Furthermore, as archaeologists grow to rely more and more on the technology, any inaccuracies in the results are less likely to be questioned, and there is a very real danger of misinterpreting the past. With the prospect of archaeologists and computer graphics specialists working even closer together in the future, there are a number of crucial issues that must be discussed to ensure the future success of such work.

This panel considers four key issues confronting archaeologists and computer graphics specialists who wish to work together to "understand the past:"

1. What is the best way to present the results? Is a significant level of realism really justified, or does this increase the chances of misleading the audience? Is a full virtual "experience" complete with virtual guide more appropriate than a simple augmented-reality headset, or even more traditional display technology?
2. What do the archaeologists really want and need from computer graphics? Are graphics practitioners listening to archaeologists, or are they simply insisting that they use the latest "cool" techniques?
3. Who pays for the work, and who benefits from the results? Is this yet another example of rich institutions exploiting resources from poorer areas for gain?
4. Will virtual tourism remove the need for anyone to visit actual sites, thus wiping out much-needed tourism revenue, or will virtual archaeology so increase the density of tourists at sites that many will have to restrict access or even close to preserve our heritage from irreparable damage?

Alan Chalmers

"Archaeology provides an excellent opportunity for computer graphics to explore scientific methods and problems at a human scale and to introduce a cultural dimension that opens up avenues for new and alternative interpretations. For example, pottery is a very common find at excavations of medieval sites. Archaeologists have been studying this material for many years, and its means of manufacture, distribution, and use are very well understood. With this sound basis provided by years of data acquisition, we can now begin to investigate less easily comprehended aspects of pottery. Medieval pottery was often very colourful, wonderfully textured, and vibrantly decorated. Was this necessary because of the lighting conditions that prevailed in medieval society, or were the pots perfectly visible and the people simply wanted some colour in their otherwise dull lives? Computer graphics offers the possibility of exploring these questions.

"I believe that in order to minimise misinterpretation of the archaeological evidence, computer visualisation of archaeological sites should be very realistic, including accurate modeling of the prevalent illumination, the complex 3D environment, and atmospheric factors such as smoke and dust. We will never know exactly why medieval pottery was so brightly coloured, but perhaps computer graphics and interactive techniques can establish a framework of possibilities within which archaeologists can hypothesise as to probable solutions."

Alan Chalmers is a senior lecturer in the Department of Computer Science at the University of Bristol, United Kingdom. He has published over 70 papers in journals and international conferences on photo-realistic graphics. He is currently vice chair of ACM SIGGRAPH. In his research, he is investigating the use of photo-realistic graphics in the visualisation of archaeological site reconstructions and developing techniques to enable virtual ancient environments to be examined under ancient illumination conditions.

Paul Debevec

"Most computer renderings of archaeological sites have done more to document the the shortcomings of graphics technology than the appearance and significance of the sites. However, a variety of new techniques such as 3D scanning and image-based rendering are making it possible to capture realistic models of complex sites and render them interactively. These developments will fundamentally change how archaeological sites and artifacts are excavated, studied, and preserved, and they will allow archaeologists to disseminate their findings in a manner more compelling than words and diagrams.

"Archaeological applications can help drive the advancement of computer graphics. The need for verisimilitude will require representing large quantities of complex geometry, recovering and rendering realistic reflectance properties, and efficient global illumination. Showing sites as they existed earlier in time will require new techniques for editing

Panelists

Kate Robson Brown
University of Bristol

Paul Debevec
USC Institute for Creative
Technologies

Zukisani Jakavula
University of Cape Town

Holly Rushmeier
IBM T. J. Watson
Research Center

acquired environments. Conveying an understanding of a people's daily lives (beyond their art and architecture) will require populating these sites with photo-realistic intelligent humans wearing photo-realistic clothing. And lastly, making complex reconstructions engaging and navigable will benefit from new forms of interactive storytelling.

"These acquisition and synthesis techniques will be expensive for several decades, but like photography and moviemaking they will become convenient and affordable. In the meantime, there is much to be gained from graphics researchers applying their work to archaeological applications; the two fields have a great deal to offer each other, and I look forward to what will come of their continued relationship."

Paul Debevec has a PhD in computer science from the University of California, Berkeley, where he is a research scientist and filmmaker. His work combines digital photography with computer graphics: deriving 3D models and reflectance properties from photographs, illuminating synthetic objects with captured light, and realistically combining real and computer generated imagery. His award-winning computer animations and art installations have featured photo-realistic renderings of the Rouen Cathedral, the University of California, Berkeley campus, and St. Peter's Basilica. His techniques have also been presented in a SIGGRAPH 2000 course (Image-Based Modeling, Rendering, and Lighting). He also recently advised on a cultural heritage preservation project in Egypt with San Francisco's Academy of Art College.

Zukisani Jakavula

"The direct equation of vision with perception appears to be implicit in almost all of the more theoretically led applications of computer graphics to the study of the archaeological record. It can be argued that a notion of perception based solely upon what can, and cannot, be seen is a gross and inadequate simplification.

"Instead, a broad definition is needed that serves to define perception not solely in visual terms, but instead as both multi-sensual reception of information about the world and the sense made of that information, mediated by factors such as memory, experience, environment, education, and expectation.

"Sight paints a picture of life, but sound, touch, taste, and smell are actually life itself. Reliance upon vision at the expense of the other senses, or perhaps more importantly, believing that there is a reliance upon the visual, can be termed 'visualism,' and it must also be acknowledged that this notion of vision as the dominant mode of perception is situated both culturally and historically within our present, Western civilisation. What has to be emphasised is that even when vision is identified as somehow dominant it never functions independently of the other senses, nor is it in any way objective. In relating vision directly with perception, computer graphics analyses must be seen as a fundamentally flawed first step rather than a clear solution to current theoretical limitations, and practitioners must once again be wary of simplification for the sake of technical convenience.

"Even when a single sensory mode is deemed primary (for example sound to the forest-dwelling Umeda of New Guinea or vision to the white middle-class academic of western Europe), it never functions in isolation. A more fully multi-sensual notion of perception, stressing the synergy of culture-nature relations provides a much fuller template upon which to proceed."

Zukisani Jakavula is a researcher in the Department of Archaeology at the University of Cape Town, South Africa. His research is investigating the use of computer graphics to correctly illuminate very realistic virtual reconstructions of prehistoric caves and shelters in the southwestern Cape, South Africa. His work considers development of novel methods that can improve perception of these ancient environments with their impressive and often enigmatic rock art.

Kate Robson Brown

"Despite the gulf of time and technology that separates them, the earliest rock art and the most sophisticated of modern computer images share a common heritage. They are artifacts of the very human effort to record the world and explore our place within it. Throughout the ages, visual images have provided a medium through which we may document, describe, summarise, order, classify, analyse, disseminate, and perhaps ultimately manage this world which we are active in both serving and creating. It is right that at this time, when developments in computer graphics have essentially made accessible a new medium for the visual image, that archaeologists and computer scientists should come together to assess the potential of this new technology for our exploration of the past in the present.



A panoramic rendering from Paul Debevec's SIGGRAPH99 film "Fiat Lux."

“Recent developments in computer graphics do not require the discipline of archaeology to construct a new theory of the visual image. Archaeologists have long been aware of the many uses and abuses of images — that they document only the intent of the image maker, that they abstract, and sometimes that they mislead. Computer graphics will benefit archaeologists with speed, accuracy, construction of massive virtual datasets, the possibility of virtual experiment, and ease of access. If there is revenue to be made out of these images, it is up to the archaeologist to manage this responsibly, honestly and professionally. This will not always be the case, but when we watch the worst happen, computer graphics per se will not be the villain of the piece.”

Kate Robson Brown is a lecturer in human origins in the Department of Archaeology at the University of Bristol, United Kingdom. She is co-founder of the Centre for Human Evolution Research, which acts as a focus for palaeoanthropological research programmes in France, Thailand, South Africa, and Zambia, and in the laboratory (hominid phylogenetics, the evolution of bipedalism). She has a strong interest in the use of computer graphics in archaeology and anthropology, and is currently involved in two projects developing this application. One involves surveying and mapping of Pleistocene decorated-cave sites in France and the Western Cape, and a second focuses on modeling the internal structure of fossil bone.

Holly Rushmeier

“The use of computer graphics in the visualization of historic artifacts offers opportunities to examine objects in novel contexts. However, computer graphics, particularly creation of 3D models, is a costly undertaking. A major issue is whether there is a way to routinely pay for this type of work, or if graphics will be relegated to a few demonstration projects.

“There are two benefits of using computer graphics: insights gained by scholars and communication of these insights to the public. We can’t put a price tag on the value of insights. To underwrite the use of computer graphics in the daily work of archaeologists and historians, we have to reduce the costs to a level within the range of grant proposals in these fields. This means not only reducing hardware and software costs, but also reducing the technical expertise required so that technical specialists are not needed. This level of cost reduction is many years off.

“Communication to the public can provide measurable income to offset the cost of technology. Admission to displays in real or virtual museums, or models delivered for use in homes or schools, can generate income. A major problem is distributing this income to sustain future research, preserve original artifacts, and cover the cost of technology. This enterprise includes the owners of the original objects, the archaeologist who develops the interpretation, the technologists who produce the displays, and finally the distributor of the final digital product. One approach is to let the market dictate — let the distributor of the digital product make the best deal possible with each of the other participants. This approach has many hazards. As a community, we need to develop guidelines for development of projects that produce income to preserve original objects, subsidize scholarship, and advance technology.”

Holly Rushmeier received a PhD degree in mechanical engineering from Cornell University in 1988. She is a research staff member at the IBM T.J. Watson Research Center. Previously she held positions at the Georgia Institute of Technology and at the National Institutes of Standards and Technology. She has served as the papers chair for the ACM SIGGRAPH conference, as papers co-chair for the IEEE Visualization conference, and as editor-in-chief of ACM Transactions on Graphics. Her research interests include data visualization, rendering algorithms, and data acquisition for image synthesis.



Hypogeum, Malta



A digital model of Michelangelo's Florentine Pieta in synthetic environment

Panelists present different angles on the latest visualization techniques coupled to the most recent research in astronomy, astrophysics, and computational fluid dynamics. The focus is on solutions to novel problems that apply to entertainment and edutainment for general audiences, movies, planetaria, and IMAX films.

John Anderson

Anderson's current interests are in understanding the use of physically motivated ideas to create a performance, cheating physical principles to span the space between real-world and "toon" physics. He came to Industrial Light + Magic in 1997 from the University of Wisconsin-Madison, where he was professor of atmospheric and oceanic sciences, associate director of the Space Science and Engineering Center, and founder of the Computational Science Program. Prior to joining the UW faculty, he was assistant professor at the University of Illinois at Urbana-Champaign and a staff engineer at MITLincoln Laboratory. He has a bachelor of science degree in physics, a masters in meteorology from the Massachusetts Institute of Technology, and a PhD in atmospheric science from Colorado State University.

Prior to joining ILM, Anderson's research was in the areas of computer simulation of physical systems, computational fluid dynamics, scientific visualization, and physically based rendering. Since joining ILM, he has worked on procedural animation for characters, which includes support for animation of clothing, fur, and skin and soft bodies. He is currently involved in creation of a set of fluid simulation tools.

Dave Bock

Dave Bock's philosophy is that the methodologies and techniques used for scientific visualization are highly dependent on the target audience and their needs or requirements. He is currently investigating high-end rendering technology for use in scientific visualization. His visualization imagery has appeared in several publications, including *Physics Today*, *Astronomy*, and *Computer Graphics World*.

Previously, he was with the North Carolina Supercomputing Center, where he developed an animation tool suite and several animated productions. These tools were showcased in the Electronic Theater at SIGGRAPH 92 and SIGGRAPH 93. He received his BS in electrical engineering at the University of Southern Illinois and teaches computer graphics at Parkland College.

woodall.ncsa.uiuc.edu/dbock/projects/VisShade/
woodall.ncsa.uiuc.edu/dbock/Vis/NeutronStar/Summary.html
woodall.ncsa.uiuc.edu/dbock/projects/AMR/

Donna Cox

Donna Cox is a pioneer in scientific visualization and has authored many book chapters, articles, and monologues on visualization in art, science, and technology. She began her pioneering work in information design in 1985 and coined the term "Renaissance teams" to describe interdisciplinary art-science collaborations. She and her work have been premiered on Good Morning America, PBS Nova, the New York Times, and the Discovery Channel. She was associate producer for scientific visualization and art director for the Pixar/NCSA segment of "Cosmic Voyage," and IMAX film that was nominated for an Academy Award in the documentary short subject category, 1997.

She was elected to the ACM SIGGRAPH Executive Committee from 1989-92 and has consistently contributed to the annual SIGGRAPH conference. She was recently appointed as editorial advisor and member of the editorial board for Leonardo, the Journal of the International Society for the Arts, Sciences and Technology.

Dennis Davidson

Dennis Davidson's current visualization interests concern understanding the environment in terms of scale, distance, surface texture, and visual patterns in living systems. Visualization subjects range from molecular and cell biology to multi-wavelength representations of the Milky Way galaxy.

He was appointed project manager for the NASA-funded Digital Galaxy Project in September 1997. He previously served for 10 years as astronomical artist at the Hayden Planetarium, where he created artwork for numerous sky shows, curated several art exhibits, and completed three major murals. He designed and implemented Hayden's computer graphics program from its modest beginning with a single Macintosh to its current array of SGI Onyx2 InfiniteReality supercomputers.

He is the museum's prime facilitator for design of Trimension's Digital Dome Display System, inaugurated at the Hayden Planetarium in early 2000 as the largest virtual reality theater in the world. His principal achievements for the Digital Galaxy Project include establishing a high-performance computer graphics studio, assembling a cross-disciplinary staff, and guiding the team towards innovative solutions for translating science data into high-definition digital movies.

He received his BA in biology from the University of California, San Diego in 1979. He is an accomplished artist and in 1992 completed a major commission for the NASA Fine Art Program. Sales of his illustrated books on space themes for children total over 250,000, while his illustrated research textbooks reach a much smaller professional audience.

Carter Emmert

Carter Emmert is dedicated to bringing awesome space visions to general audiences. He tries to be as true to data as possible, though not all data can be verified.

He has been a scientific illustrator and visualizer for 20 years, working with NASA, the Smithsonian Air and Space Museum, and NCAR. He has a diverse background in conceptualization and visualization of future space exploration, Earth-system simulation, and astronomical phenomena. He is currently directing the visualization efforts of the Digital Galaxy Project, where models of the cosmos are turned into images that will fill a 68-foot dome at the newly rebuilt Hayden Planetarium in New York using developmental technology and techniques.

Together with the Hayden planetarium and the San Diego Supercomputer Center, he has modeled the clouds of dust and gas in the Orion nebula using volumetric techniques. Traditional computer graphics techniques model reality using surfaces built from polygons and splines. Cosmological objects, are rarely so simply defined. For instance, the Orion nebula is an immense cloud of dust and gas two dozen light-years across and 1500 light-years away.

Robert Patterson

Robert Patterson's approach is to search a simulation for the most interesting aspects and highlight those aspects with a cinematic and artistic approach. His rendering methods often augment or exaggerate reality to "see" the data, which might be dark matter 500 million light years away. He uses immersive virtual reality as a virtual design workspace for creating cinematic presentations of scientific visualization. He used this system to create camera choreographies for the I MAX film "Cosmic Voyage," the Hayden Planetarium grand opening, and the upcoming PBS NOVA program: "Mapping the Universe."

He has collaborated on visualizations at NCSA since 1989, except for a year at ILM in 1994 to work on "Forrest Gump." He is currently teaching a course for the New Media Program at Indiana University, Indianapolis. He received Bachelor of Science degrees in film production and advertising from the University of Texas at Austin.

virdir.ncsa.uiuc.edu/virdir/universe/
virdir.ncsa.uiuc.edu/virdir/raw-material/cosmology/AMR/
virdir.ncsa.uiuc.edu/virdir/raw-material/general/VR/

Dennis Davidson
American Museum
of Natural History

Carter Emmert
American Museum
of Natural History

Robert Patterson
National Center for
Supercomputing
Applications

Stephen Rosenbaum
Industrial Light + Magic

Michael Wahrman
New York University

Stephen Rosenbaum

“Creating imagery for movies that simulate some physically real event often leads to bending scientific accuracy to suit the story. Computer graphic artists frequently struggle with generating a dynamically correct representation of the effect while trying to accommodate the director’s desire to exaggerate the visuals in order to stimulate some profound feeling from the audience. In short, if it’s technically accurate but boring to look at, it won’t work for the movie. We call this process taking creative license.”

Stephen Rosenbaum joined Industrial Light + Magic in 1988 as one of the original members of the newly formed computer graphics department. He played an integral role as an artist and supervisor on several of the most groundbreaking digital effects films of this decade, including five that won Oscars for best achievement in visual effects. He received an Academy Award and a British Academy Award for his work on “Forrest Gump.” In 1996, he left ILM to continue working with his mentor, Ken Ralston, who was training him to be a visual effects supervisor. He returned to ILM in 1998. He is a member of the Academy of Motion Picture Arts and Sciences and the Association for Computing Machinery, and a founding member of the Visual Effects Society. He studied computer graphics, film, photography, and theater lighting and scenic design at the University of California, Berkeley, where he received his bachelors degree in 1987.

Michael Wahrman

Michael Wahrman’s position has evolved over time. In the past, as a designer of rendering solutions for the motion picture industry, he believed that one could achieve a richer and more inherently realistic image by understanding the physical processes that underly natural phenomenon and incorporating them into production methodologies. Since working as an advisor on scientific visualization projects, he has come to believe that the worlds of computer animation and visualization are nearly totally different.

He is a visual effects supervisor and consultant who divides his time between the motion picture industry and research. His professional experience in production began when he left the Rand Corporation in 1981 to design and implement a new animation system at Robert Abel and Associates, a pioneering production studio. Since then he has worked either in a leadership role or as a consultant in computer animation and visual effects projects. His clients have included Universal Studios, Walt Disney Studios, Paramount Landmark Entertainment, Sony, IBM Research, Manex Entertainment, and many others. Recent projects have included “Starship Troopers” and “What Dreams May Come.” His films have received awards from the Chicago Film Festival, Imagina, Ars Electronica, and other organizations. He is a four-time member of the jury of Ars Electronica, two-term member of the board of directors of the New York Professional Chapter of ACM SIGGRAPH, and a member of the Visual Effects Society. In 1997, he received an Academy of Motion Picture Arts and Sciences Technical Achievement Award. He is currently a senior research scientist at New York University’s Media Research Laboratory. In 1998, he took a year off from motion picture work to become visual effects advisor for the American Museum of Natural History at the Hayden Planetarium.

A new SIGGRAPH program, Art & Culture Papers, is launched at this Special Session. These three papers present cutting-edge questions about the role of computer graphics and interactive techniques in today's art and culture. Speakers consider the changing roles of technical tools (artificial intelligence practices), popular cultural forms (science fiction), and longstanding philosophical issues (the distinction between an idea and its representation) in a world of computer visualization, biotechnology, net.art, and the melding of art and technology for cultural production. The papers are published in their entirety in the SIGGRAPH 2000 Electronic Art & Animation Catalog.

DIGITAL ONTOLOGIES: THE IDEALITY OF FORM
IN/AND CODE STORAGE — OR — CAN GRAPHESIS
CHALLENGE MATHESIS?

Johanna Drucker

Abstract: The attempt to understand the connections that link human thought to its representation through the act of formgiving (in language, image, or signs) is central to Western philosophy and aesthetics. This paper examines the way such issues, particularly the question of whether an idea can exist outside of instantiation in material form, are newly posed with respect to the digital environment. Does the commonality of code storage as the defining condition of digital processing confirm a long-standing Western philosophical quest for mathesis (in which knowledge and its representation exists in a perfect, unambiguous form)? Can this claim to absolute, clear, "truth" be challenged by graphesis (the constitution of information within visual, graphic modes)?

Johanna Drucker is the Robertson professor of media studies at the University of Virginia. She has published and lectured extensively on topics related to the history of visual communication, typography, writing, and visual art. Her scholarly books include: "Theorizing Modernism" (Columbia University Press, 1994), "The Visible Word: Experimental Typography and Modern Art" (University of Chicago Press, 1994); "The Alphanumeric Labyrinth: The Letters in History and Imagination" (Thames and Hudson, 1995); and "The Century of Artists' Books" (Granary, 1995). She has been printing and publishing creative work since the early 1970s and has produced more than two dozen books, many using experimental typography, through her own press, Druckwerk. Recent titles include "Simulant Portrait" and "Otherspace: Martian Typography" (a collaboration with artist Brad Freeman), both fictional works dealing with futuristic identity and communication. She has written and lectured on the subject of language in the electronic environment, the status of art in the age of information, and various philosophical issues raised by virtual media, and is currently at work on a book on art as information.

EXPRESSIVE AI

Michael Mateas

Abstract: The field of artificial intelligence (AI) has produced a rich set of technical practices and interpretive conventions for building machines whose behavior can be narrated as intelligent activity. Artists have begun to incorporate AI practices into cultural production (into the production of artifacts and experiences that function as art within the cultural field). I term my own practice of AI-based cultural production Expressive AI. I will attempt to provide a preliminary understanding of this practice by both situating expressive AI with respect to other discourses on AI and by working inductively from my own AI-based art work. I will first provide a brief description of three of my AI-based art pieces. These will serve as concrete examples to ground the rest of the discussion. I will then describe the expressive AI practice by first situating it with respect to the GOFAI/interactionist AI debate, then by describing the central organizing metaphors of authorial and interpretive affordance, and finally by providing a preliminary set of desiderata for expressive AI practice.

Michael Mateas' work explores the intersection between art and artificial intelligence. He received his BS in engineering physics from the University of the Pacific and his MS in computer science from Portland State University. He is currently a PhD student in computer science at Carnegie Mellon University and a research fellow in the art department's Studio for Creative Inquiry. His artwork includes: Office Plant #1, a sculptural installation that responds physically to the content of email being received at the computer terminal where it is "planted;" Terminal Time, a machine that constructs ideologically-biased documentary histories in response to audience feedback; and Subjective Avatars, virtual world avatars that actively filter a user's experience of a virtual world so the user inhabits a specific subjective position. His work has been exhibited in the Digital Salon in New York, the Carnegie Museum, the Warhol Museum, and at numerous conferences and festivals. Before he began his work at Carnegie Mellon, he worked at Intel Laboratories, where he co-founded GEAR (Garage Ethnography and Applications Research), a research group that employs ethnographic techniques to understand how new computing technology fits into people's lives.

Speakers
Johanna Drucker
University of Virginia

Michael Mateas
Carnegie Mellon University

Eugene Thacker
Rutgers University

THE SF OF TECHNOSCIENCE: THE POLITICS OF
SIMULATION & A CHALLENGE FOR NEW MEDIA ART

Eugene Thacker

Abstract: Developments in biotechnology are at the forefront of what the future of medicine, health care, and the body will come to mean in technologically advanced cultures. This paper considers the tensions and inconsistencies within the biotech industry as a zone where science fiction (SF) is put to work as negotiator and mediator. Taken as a discursive practice, SF constructs futuristic narratives of progress and conditions the very techniques, technologies, and research that may take place. However, as cultural theorists such as Donna Haraway and Fredric Jameson suggest, SF can also fulfill a critical function, highlighting the contingencies and limitations in biotech's self-fulfilling narrative of future medicine. A consideration of the emerging category of "net.art" provides one starting point for a critical SF as practice.

Eugene Thacker teaches in the Program in Comparative Literature at Rutgers University, where he directs [techn] New Media & Digital Arts. He is a contributing editor at *The Thing* and a collaborator with *Fakeshop*.

Diane Gromala

Diane Gromala is associate professor in the School of Literature, Communication, and Culture at Georgia Tech, where she teaches in the graduate program in Information Design and Technology. She is also a faculty member of the Graphics Visualization Center. Her courses focus on the cultural and theoretical aspects of new and virtual media. Gromala's critical analyses of virtual media are informed by her work as an artist and designer in virtual reality and other forms of digital media. Her articles have appeared in numerous art and design journals, and her virtual reality artwork has been performed and presented in Canada, the US, Europe, and the Middle East, and aired on the Discovery Channel and the BBC. She is on the Editorial Board of *Postmodern Culture*, Chair of the SIGGRAPH 2000 Art Gallery, and a recent recipient of a Senior Fulbright Fellowship.

Noah Wardrip-Fruin

Noah Wardrip-Fruin is an artist, writer, and research scientist at the New York University Media Research Lab. His current projects include *The Impermanence Agent*, a storytelling web agent that customizes its story of impermanence for each user until none of the original story remains. *The Impermanence Agent* is featured in the SIGGRAPH 2000 Art Gallery. His previous work includes *Gray Matters*, a collaborative fiction embedded within images of the human body, which he completed as an Edward Albee Foundation Fellow. Its exhibition at the Sandra Gering Gallery was the first public installation of a zooming user interface.

What is the future of fiction? Members of the graphics, interaction, literature, and arts communities answer that question and imagine the shape of fiction in the emerging realms of networked computers. Following up on SIGGRAPH 99's highly successful session, new speakers engage in a roundtable discussion with last year's participants.

Espen Aarseth
"Forget narrative!"

"In designing interesting fictional worlds (games, didactic tools etc.), we rely far too much (to the point of fetish) on the structures of storytelling and narrative. As with any new medium, it takes time (50 years?) to develop sound, autonomous aesthetics that exploit a new technology fully, and the sooner we stop thinking in terms of story and plot, and start to take the user seriously as a co-creator of meaningful content, the better."

Jay Bolter

"We cannot know the future of fiction, but we can talk about its present condition. What strikes me is the number and variety of narrative media forms today: from traditional novels and films to Web sites and interactive games. There is no one form around which our culture's creative energies are converging."

Andrew Glassner

"The phrase "interactive fiction" is the progeny of a shotgun marriage between two innocent words. Despite millennia of experimentation, direct audience participation in stories has never taken off, because the classical story form requires authorial control. Exciting new forms of storytelling will use computers and networks to provide a structured environment for guided, cooperative imagination."

Michael Joyce

"The future of fiction is in its past and ours. Fiction, in whatever form, calls us to find purpose in surface, to find a lasting world in our shifting sense of ordinary life. It is a call to ordinary mindfulness. Surface attentions grace human life and grant us the peculiar joys that actualize even increasingly virtual lives."

Jesse Schell

"Story and gameplay are two different sides of the same coin."

Phoebe Sengers

"The future of fiction will be split in two. The mega-corporations that are vying for domination of the Web will inundate it with globalized 3D interactive repeats of tired narratives of the late 20th century ("Gentlemen Prefer Blondes," now with interactive Marilyn!). Insurgent writers and artists will fight back with atomized, decentralized, emergent post-narratives, elegant in their construction, intellectually dense, extremely hip, but not exactly a pleasure to read. A possible third alternative: technology that supports our need and desire to understand the world and express ourselves narratively."

Ana Serrano

"The digital story artist is not just a griot with Flash, but someone charged with a task demanding more generosity and imagination: not the "telling" of stories, per se, in this new medium, but creation of a singular, compelling experience that melds the activities of story 'viewing-creating-sharing.' She blazes a preferred set of paths that each traveler navigates in his own fashion, with the opportunity to add side-trails, scenic lookouts — even trolls — for travelers who happen along that way later."

Adrienne Wortzel

"If, in the act of seeing, we are always subject to being seen, then a mirror will always be nearby in every direction in our mind's eye, reflecting what we think may be visible to others. This turns us into actors and renders the world a more theatrical place, a global amphitheater where the border between the occurrence and portrayal of events has dissolved. The threshold of fiction is now a proscenium enabled by new technologies."

Andrew Glassner
Consultant

Michael Joyce
Vassar College

Jesse Schell
Walt Disney
Imagineering

Phoebe Sengers
GMD National Research
Center for Information
Technology

Ana Serrano
Canadian Film Centre

Adrienne Wortzel
City University
of New York

Espen Aarseth

is associate professor of humanistic informatics at the University of Bergen, and the author of numerous essays on hypertext aesthetics, online culture, and computer games, as well as the book "Cybertext: Perspectives on Ergodic Literature" (Johns Hopkins, 1997, www.hf.uib.no/cybertext/). He directs the CALLMOO project (cmc.uib.no), which uses MOOs in foreign language learning. Currently, he is engaged in writing a book on digital culture and politics.

Jay Bolter

is director of the New Media Center and Wesley Professor of New Media at the Georgia Institute of Technology. Together with Michael Joyce, he is the author of Storyspace, a program for creating hypertexts for individual use and World Wide Web publication. He is the author of "Turing's Man" (1984), "Writing Space" (1991), and most recently, "Remediation" (1999), coauthored with Richard Grusin.

Andrew Glassner

is a writer, researcher, consultant, and speaker in the fields of computer graphics, story structure, and storytelling. He started working in graphics research in 1978 at the New York Institute of Technology Computer Graphics Lab, and has carried out work at Case Western Reserve University, the IBM T.J. Watson Research Lab, the Delft University of Technology, Bell Communications Research, Xerox PARC, and Microsoft Research. He is the author of a regular column for IEEE Computer Graphics and Applications and several books, including *3D Computer Graphics: A Handbook for Artists and Designers*, *Principles of Digital Image Synthesis*, and most recently, *Andrew Glassner's Notebook*, a compilation and extension of his IEEE Computer Graphics and Applications columns. He is also creator and editor of the Graphics Gems series and the book *An Introduction to Ray Tracing*. He has served as Papers Chair for SIGGRAPH 94, founding editor of the *Journal of Graphics Tools*, member of the Editorial Board of *IEEE Computer Graphics and Applications*, and editor-in-chief of *ACM Transactions on Graphics*. He directed the short film "Chicken Crossing," which premiered at the SIGGRAPH 96 Electronic Theater, and designed the highly participatory game "Dead Air" for the Microsoft Network, where he wrote and directed the live-action pilot episode. He holds a PhD from the University of North Carolina at Chapel Hill. His consulting practice specializes in the creative use of 3D graphics, and in techniques for combining story and storytelling in games and online entertainments. He is currently at work on his second novel.

Michael Joyce

is perhaps best known as the originator of an art form. The New York Times called his "afternoon" "the granddaddy of hypertext fictions," while The Toronto Globe and Mail said that it "is to the hypertext ... interactive novel what the Gutenberg Bible is to publishing," and Der TAZ has called him "Der Homer der Hypertexte." His hyperfictions, "Twilight, A Symphony" on CD-ROM and "Twelve Blue" on the World Wide Web, were published by Eastgate, as was the most recent "Lasting Image," with Carolyn Guyer (2000). Another Web fiction, "On the birthday of the stranger..." was published in the inaugural electronic issue of the Evergreen Review. A linear novel, "Going the Distance," is published on the Web by Pilgrim Press. His new collection of essays, "Othermindedness: the emergence of network culture," was published in January 2000 by the University of Michigan Press, which also published his essays on hypertext theory and pedagogy: "Of Two Minds: Hypertext Pedagogy" and "Poetics" (1995). He serves on the editorial boards for PostModernCulture, Works& Days, and the Computers and Composition Journal. He is currently associate professor of English and director of the Center for Electronic Learning and Teaching at Vassar College.

Jesse Schell

is a show programmer and game designer at the Walt Disney Imagineering Virtual Reality Studio. He has helped develop such attractions as: "Aladdin's Magic Carpet Ride," "Hercules in the Underworld," "Mickey's Toontown Tag," and an interactive "Pirates of the Caribbean," all for DisneyQuest (Disney's chain of interactive indoor theme parks) or Disneyworld. His pre-Disney work includes: designing artificial intelligence systems for automated storytelling as part of the Rensselaer Institute's Autopoiesis project; designing and building the NVR system for networked virtual reality at Carnegie Mellon University (an integral part of the Virtual Reality: An Emerging Medium exhibit at the Guggenheim Soho in New York City); and writing, directing, and performing as a juggler, comedian, and circus artist in shows with both the Juggler's Guild and Freihofer's Mime Circus entertainment troupes.

Phoebe Sengers

is a research scientist in Media Arts Research Studies at the GMD National Research Center for Information Technology. She graduated in 1998 from Carnegie Mellon University with a PhD in artificial intelligence and cultural theory. In her thesis, she developed an agent architecture that allows agents to structure their behavior so that it is narratively understandable. In 1998-99, Sengers was a Fulbright Guest Professor at the Center for Art and Media Technology (ZKM) in Karlsruhe, Germany. She co-chaired the 1999 AAI Fall Symposium on Narrative Intelligence.

Ana Serrano

is the director of Bell h@bitat, the new media training facility at the Canadian Film Centre, a world-renowned film, television, and new media institute established by Norman Jewison. She oversees the strategic planning, program design, and fiscal development of all of the centre's new media initiatives, including creation of interactive narrative prototypes through the centre's New Media Design Program. Featured as one of MacLean Magazine's Top 100 Canadians To Watch in the year 2000, she is currently producing the Great Canadian Story Engine Project, an online oral history Web site of personal Canadian stories, developed in the Canadian Film Centre's New Media Design Program. She is active on the boards of the Canadian Conference on the Arts, the Information Technology Design Centre at the School of Architecture at the University of Toronto, and the Muriel Cooper Prize Council at the Design Management Institute. She frequently speaks at new media and film festivals throughout the world about the emerging realm of interactive and networked digital storytelling.

Adrienne Wortzel

is a new media artist and educator. Her artworks explore the range of new technologies as Web works, robotic installations, and online and offline performance productions and writings. Her Globe Theater Repertory Company of robots were featured in installations at Technoseduction-Cooper Union; Creative Time's Art in the Anchorage at the Brooklyn Bridge; Ars Electronica, Linz, Austria; and Kunstpflug e.V., Berlin. She is the recipient of an artist-in-residence grant at Lehman College Art Gallery in New York funded by the Electronic Media and Film Program of the New York State Council on the Arts, where she produced "Sayonara iorama," a full theatrical performance with robots, live performers, and responsive remote performances via videoconferencing technology. Recently, she was one of two jurors for the first New York

University NYU Press Prize for electronic hyperfiction. She is associate professor of advertising design and graphic arts at New York City Technical College of the City University of New York, as well as visiting professor at the Cooper Union for the Advancement of Science and Art as a result of a National Science Foundation grant awarded for creating robots as actors for theater for the year 2000.

Diane Gromala

is associate professor in the School of Literature, Communication, and Culture at Georgia Tech, where she teaches in the graduate program in Information Design and Technology. She is also a faculty member of the Graphics Visualization Center. Her courses focus on the cultural and theoretical aspects of new and virtual media. Her critical analyses of virtual media are informed by her work as an artist and designer in virtual reality and other forms of digital media. Her articles have appeared in numerous art and design journals, and her virtual reality artwork has been performed and presented in Canada, the US, Europe, and the Middle East, and aired on the Discovery Channel and the BBC. Gromala is on the editorial board of Postmodern Culture, Chair of the SIGGRAPH 2000 Art Gallery, and a recent recipient of a Senior Fulbright Fellowship.

Noah Wardrip-Fruin

Noah Wardrip-Fruin is a fiction writer, artist, and research scientist at the New York University Media Research Lab. His current projects include The Impermanence Agent, a storytelling Web agent that customizes its story of impermanence for each user until none of the original story remains. The Impermanence Agent is featured in the SIGGRAPH 2000 Art Gallery. His previous work includes Gray Matters, a collaborative fiction embedded within images of the human body, which he completed as an Edward Albee Foundation Fellow. Its exhibition at the Sandra Gering Gallery was the first public installation of a zooming user interface.



Sayonara Diorama, Adrienne Wortzel, Darwin, Pandora and the Mappa Mundi Monsters, 1998

This Special Session features individuals who played key roles in the production of ground-breaking digital effects for the feature film, “The Perfect Storm.” Each presenter delivers a 10-15 minute overview that focuses on one of the unique challenges involved in simulating realistic ocean waters and the drama that subsequently unfolds.

Tim Alexander, Digital Compositor Supervisor

Summarizes bluescreen extractions with mist and water, in addition to depth compositing for adding haze and depth of field effects.

Tim Alexander graduated in 1994 with a masters degree in electrical engineering from Cornell University, where he was studied lighting design in the Advanced Undergraduate Theatre Program. He worked for Disney for three years, including a year-long internship while he was still a student. While at Disney, he worked on “James and the Giant Peach.” He joined Industrial Light + Magic in 1996.

ILM Credits

- 2000 “The Perfect Storm” - Compositor
- 1999 “Deep Blue Sea” - Compositor
- 1999 Star Wars: Episode I “The Phantom Menace” - Compositor
- 1997 “Titanic” - Compositor
- 1997 “The Lost World: Jurassic Park” - Compositor
- 1996 “Star Trek: First Contact” - Compositor

John Anderson, CG Scientist

Talks about the physical simulation processes used for the animation of water. His discussion includes an overview of the techniques used for generation of ocean wave fields, high-resolution simulations for hero wave dynamics, and techniques for treating flows, which interact with boat geometry. He came to ILM in 1997 from the University of Wisconsin-Madison, where he was professor of atmospheric and oceanic sciences, associate director of the Space Science and Engineering Center, and founder of the Computational Science Program. Prior to joining the UW faculty, he was assistant professor at the University of Illinois at Urbana-Champaign and a staff engineer at MIT Lincoln Laboratory. He has a bachelor of science degree in physics, a masters in meteorology from the Massachusetts Institute of Technology, and a PhD in atmospheric science from Colorado State University.

Prior to joining ILM, Anderson’s research was in the areas of computer simulation of physical systems, computational fluid dynamics, scientific visualization, and physically based rendering. Since joining ILM, he has worked on procedural animation for characters, which includes support for animation of clothing, fur, and skin and soft bodies. He is currently involved in creation of a set of fluid simulation tools.

ILM Credits

- 2000 “The Perfect Storm” - Research and Development
- 2000 “Mission to Mars” (in production) - Research and Development
- 1999 Star Wars: Episode I “The Phantom Menace”
Research and Development
- 1999 “The Mummy” - Research and Development
- 1998 “Mighty Joe Young” - Software Development

Stefan Fangmeier, Visual Effects Supervisor

Introduces the film and talks about some of the goals and new techniques developed for this production. He joined the computer graphics department at ILM in 1990. His first major project was “Terminator 2: Judgment Day.” Since then, he has worked on various projects, including “Hook” and “Jurassic Park” as a computer graphics supervisor and on “Casper” as the digital character co-supervisor. In 1997, Fangmeier won a British Academy Award for his work as visual effects supervisor on “Twister.”

Prior to joining ILM, he served as director of production at Mental Images GmbH and Co. in Berlin. In addition, he has worked as a scientific visualization program manager at the National Center for Supercomputing Applications at the University of Illinois, as a technical director for Digital Productions in Los Angeles, as a computer operator, and as a systems programmer and image processing analyst at the Aerospace Corporation in California. He received his degree in computer science from California State University, Dominguez Hills in 1983.

ILM Credits

- 2000 “The Perfect Storm” - Visual Effects Supervisor
- 1999 “Galaxy Quest” - Visual Effects Consultant & Second Unit Director
- 1998 “Saving Private Ryan” - Visual Effects Supervisor
Recipient of British Academy Award
Best Achievement in Visual Effects
- 1998 “Small Soldiers” - Visual Effects Supervisor
- 1997 “Speed 2: Cruise Control” - Visual Effects Supervisor
- 1996 “Twister” - Visual Effects Supervisor
Recipient of British Academy Award
Best Achievement in Visual Effects
Nominee for Academy Award
Best Achievement in Visual Effects
- 1995 “Casper” - Digital Character Co-Supervisor
- 1993 “Jurassic Park” - Lead Computer Graphics Supervisor
- 1991 “Hook” - Computer Graphics Supervisor
- 1990 “Terminator 2: Judgment Day” -
Computer Graphics Shot Supervisor

Doug Smythe, Associate FX Supervisor

Addresses the issues and special concerns of lighting, shading, and rendering digital oceans. He also discusses some of the methods used in the pipeline between the simulation and the rendering phases.

Doug Smythe joined Industrial Light + Magic in June of 1987 after receiving his bachelor of science degree in electrical engineering and computer science from the University of California, Berkeley. He has received three awards from the Academy of Motion Picture Arts and Sciences: a Technical Achievement Award for invention of the “morph” process, a Scientific and Engineering Award for co-invention of a digital motion picture retouching system, and an Oscar for Best Achievement in Visual Effects for his work in “Death Becomes Her.” He has been interviewed for dozens of magazine articles and television shows, and has given several lectures at the University of California, Berkeley and various international conferences, including SIGGRAPH and Imagina.

ILM Credits

- 2000 “The Perfect Storm” - Associate Visual Effects Supervisor
- 1999 Star Wars: Episode I “The Phantom Menace”
Computer Graphics Supervisor
- 1996 “101 Dalmations” - Computer Graphics Supervisor
- 1995 “Jumanji” - Computer Graphics Sequence Supervisor
- 1994 “Star Trek Generations” -
Computer Graphics Technical Supervisor
- 1994 “Baby’s Day Out” - Computer Graphics Supervisor
- 1993 “Meteor Man” - Computer Graphics Technical Director
- 1992 “Jurassic Park” (Trailer) - Computer Graphics
Supervisor/Technical Director
- 1992 “Death Becomes Her” - Computer Graphics Supervisor
Academy Award - Best Visual Effects
British Academy Award - Best Visual Effects
- 1992 “Memoirs of an Invisible Man” -
CG Supervisor/Technical Director
- 1991 “Terminator 2: Judgment Day” -
CG Supervisor/Technical Director

Habib Zargarpour, Associate FX Supervisor

Shows how John Anderson’s ocean simulations were used in Maya, how the boats were setup to be simulated over them, and how “targetCams” were mounted on another boat to allow ocean scouts to compose some shots. He also discusses some of the splash spray and mist simulation techniques.

Habib Zargarpour joined ILM as a technical director on “The Mask” after working as an illustrator since 1981. He received his BASC in mechanical engineering from the University of British Columbia and went on to graduate with distinction in industrial design from the Art Center College of Design. He stumbled into 3D graphics in 1990 while designing for a film, and he considers computer graphics to be “the perfect mix” of the technical and artistic worlds. He still likes to paint with real brushes to maintain his sanity.

ILM Credits

- 2000 “The Perfect Storm” - Associate Visual Effects Supervisor
- 1999 Star Wars: Episode I “The Phantom Menace” -
Computer Graphics Supervisor
- 1997 “Spawn” - Associate Visual Effects Supervisor
- 1996 “Star Trek: First Contact” - Computer Graphics Supervisor
- 1996 “Twister” - Digital Tornado Designer/
Computer Graphics Supervisor
Recipient of British Academy Award
Best Achievement in Visual Effects
Nominee for Academy Award
Best Achievement in Visual Effects
- 1995 “Jumanji” - Computer Graphics Sequence Supervisor
- 1994 “Star Trek: Generations” - Senior Technical Director
- 1993 “The Mask” - Technical Director

Phil Tippett's History of Animation

Phil Tippett presents a history of visual effects from his unique perspective, based on his extensive experience developing creature and characters for feature films. From his beginnings at Cascade Studios, where he animated the Pillsbury Doughboy, to becoming the head of a full-service digital animation studio, he traces the development of animation from stop-motion to computer graphics. The Academy Award-winning VFX supervisor also discusses the influences of other artists on his work and the challenges of running a digital studio in the current marketplace.

Phil Tippett

Phil Tippett, founder of Berkeley-based Tippett Studio, is a director of animation and producer of visual effects. In 1976, he joined director/producer George Lucas to animate the miniature holographic chess match between R2-D2 and Chewbacca, and create characters for the famed Cantina scene in "Star Wars." Two years later, he returned to Northern California to head George Lucas' ILM animation department with Jon Berg for the Star Wars sequel, "The Empire Strikes Back." In 1982, he co-developed the animation technique called Go-Motion for "Dragonslayer" and received his first Academy Award nomination.

By 1983, Tippett was head of the famed Lucasfilm "creature shop" and was awarded his first Oscar, for "Return of the Jedi" in 1984. After "Return of the Jedi," Phil left ILM and set up a studio in his garage to create a 10-minute experimental film called "Prehistoric Beast," which led to the 1985 CBS animated documentary, "Dinosaur." Tippett Studio won its first award, an Emmy for Special Visual Effects, for the animated dinosaur sequences for this CBS show.

In 1985, Tippett created the animated robot sequences for "Robocop," beginning a long collaboration with Paul Verhoeven.

In 1991, Steven Spielberg selected Tippett to supervise animation on 50 dinosaur shots for "Jurassic Park." As animation director, he supervised both the Tippett Studio and the ILM animators. The resulting realistic portraits of dinosaurs that breathe, flex, twitch, and react earned a second Oscar.

With the transition from stop-motion animation to computer graphics, Tippett says: "The myth in this business is that now special effects and dimensional character animation is easier to do. All you have to do is push the 'Better' button on the computer and somehow the look of a shot improves. But actually it is our understanding of the traditional crafts like performance, picture making, and sculpting that drives the work we do with these new tools."

In 1995, Tippett Studio created the giant, hostile alien arachnids in "Starship Troopers," released in November of 1997. For this massive endeavor, Tippett Studio produced over 250 creature effects shots and created an all-CGI facility. The excellence of the work in this film resulted in Tippett's sixth nomination for an Academy Award.

The late 1990's have been a time of substantial growth at Tippett Studio. Tippett has supervised animation and effects in "Virus" for Universal and "My Favorite Martian" for Disney, and recently completed production on "The Haunting" and "Komodo." He has recently completed co-supervising (with Craig Hayes) the effects on "The Haunting" for DreamWorks SKG and is developing his own animated features as well as other live-action/VFX theatrical features. Tippett Studio is currently in production on Paul Verhoeven's "The Hollow Man."



Phil Tippett on the set of DreamWorks' "The Haunting."
Photo by Frank Masi

We are all standing on the precipice of the chasm that stretches from today to the future, where the walls surrounding what is perceived as reality begin to blur with the synthetic worlds we create.

This is a truly exciting age. For the first time, we are at the threshold where the technology has caught up with developers' seemingly insatiable need for more image-generation power. Nevertheless, a whole new set of problems is now looming on the horizon. How do we harness this expressive power? How do we deliver interactive impact over the Web? How do we make 3D graphics so intuitive that anyone can use them?

The Web3D RoundUP is a high-speed shootout that showcases the latest and very greatest interactive 3D content available over the Web. Top 3D tools and content developers demonstrate their jury-selected Web3D technology and applications in a fast-paced, exciting format. Diplomatically speaking, we describe this as an "interactive event for interactive content." Realistically, it's a monster. Have you ever tried doing a cool demo in less than three minutes, with the seconds ticking down on the big screen, in front of 1700+ world-class experts armed with happy/sad sound effects and Nerf guns?

Twice a year, at the annual SIGGRAPH conference and at the Web3D/VRML Symposia, Web3D RoundUP presents the best the Web3D world has to offer in a cool and entertaining way. The challenge for

jurors is to choose two dozen diverse, eyepopping selections from eight dozen submissions, two weeks before SIGGRAPH 2000! The challenge for presenters is to engage the audience in a fun way, so they can interactively participate by using the various "interactive feedback devices" and become part of the event itself. The challenge for the audience is to hold on during this wild ride and then vote to pick the best entries in each category.

Our community's future lies partially within the infrastructure that the online gaming developers and console game manufacturers are creating for us. It still boggles the mind to think that we can now buy cheap little black boxes that surpass the performance of a \$100,000 super-computer of only a few years ago. As an industry, and as individuals, are we ready to harness this power? Everyone has been waiting 10 years for this moment, where the Internet and real-time 3D begin to intersect. Now they are slamming together like particles in a supercollider. These impacts are sparking an entirely new breed of applications, content, and experiences, unlike any the world has ever experienced, because they can be easily shared.

The SIGGRAPH 2000 event is the 10th Web3D RoundUP. It evolved from the early days of the Virtual Reality Education Foundation, affectionately known as VeRGe, which was created by Timothy Childs, Linda Jacobson, and Peter Rothman. It was at those early monthly VeRGe meetings (and the early SIGGRAPH Demo SIGS organized by Don



Brutzman and populated by the VRML community) that we got our first tastes of technically successful chaos, producing and presenting bleeding-edge demo sessions that morphed into today's Web3D RoundUP. This year's Web3D RoundUP sets up an inspirational face-off between "regular" Web3D computer graphics and the real-time gaming communities.

One of the hot spices that gave the Web3D RoundUP its original kick was the oft-occurring computer crash. There seemed to be a direct correlation between a crashing computer and how bleeding-edge the demo was. Now Web3D technology is starting to stabilize (and fewer fire extinguishers are needed), but we still seek the bleeding edge. Where is it? A big chunk lies in the gaming community, where programmers have moved past platform stability and are harnessing the new power. It is imperative that dialogue occurs in many different forms between these two worlds. There are huge lessons that game designers can learn from the computer graphics world and vice versa, especially in understanding practical human-factors issues in immersive environments.

The Web3D RoundUP became what it is today thanks to great help from all the volunteers at the VRML/Web3D and SIGGRAPH conferences, as well as the solid advice and support by many SIGGRAPH chairs, staff members, and committee members. We are most fortunate to receive sage guidance and adult supervision from the world's greatest audiovisual wizards at AVW. Finally, we thank the hard-working technical trapezists and content creators, who put in countless hours getting ready for just a few minutes of intense glory onstage.

Time's up, gotta go. Boom! Next!



Contents

162 Benjamin Franklin High School

166 Bateaux2000

166 RVV (Remote Visual Vehicle)

SIGKIDS

MARC J. BARR
MIDDLE TENNESSEE STATE UNIVERSITY



Committee

CHAIR

Marc J. Barr
Middle Tennessee State University

COMMITTEE

Marc J. Barr
Middle Tennessee State University

Mk Haley
Walt Disney Imagineering

Scott Lang
The Academy for the Advancement of Science and
Technology

Paul Werner
Benjamin Franklin High School

Sarah Roberts
Louisiana Children's Museum

Erica Hornung
Freelance Digital Skylarking

Guanping Zheng
Middle Tennessee State University



Introduction

THIS YEAR, sigKIDS FOCUSED ON GETTING INVOLVED with the local community and branching out to a new audience.

In the spring of 1999, discussions were begun among sigKIDS, Community Outreach, representatives of the local schools, and the Louisiana Children's Museum in New Orleans. During the fall of 1999, a group of students from Benjamin Franklin High School, a public magnet school in New Orleans, led by their teacher, Paul Werner, began to work on a Web site for sigKIDS. At SIGGRAPH 2000, they worked with SIGGRAPH TV/Online to produce Web content. The Web site also contains artwork and animations submitted by students and teachers from the US and Europe. This Web site will remain online until July 2001.

SIGGRAPH 2000 has arranged for the Louisiana Children's Museum to receive 15 computers donated by a generous local corporation. With this technology, the museum developed its first hands-on lab for visitors.



Benjamin Franklin High School, a New Orleans public magnet school, was chosen to produce the Web site and online video for sigKIDS 2000. Paul Werner, Franklin's Webmaster and television production teacher, headed the team composed of students from 9th through 12th grades. All interested students were invited to join the team and learn the skills necessary to produce the material for the conference.

STUDENT COMMENTS

Leslie Becker

"I joined the SIGGRAPH 2000 project in September of 1999 when Marc Barr came to visit Franklin. I thought it would be a wonderful way for me to learn basic ideas on advanced computer programs. In November 1999, I attended the SIGGRAPH 2000 educational outreach program at Kennedy High School. In February 2000, I attended and interviewed Marty Sixkiller at the New Orleans Museum of Art. A brief summary of Marty Sixkiller's life and the interview can be found on the egroup site. I was able to watch an educational video and learn how the advanced crew and technology at Pacific Data Images conquered some graphic complications and created great computer-animated films such as "ANTZ." This summer, I am going to India to film/direct/produce an educational video about prenatal care in India. I will be shooting with a Cannon G1-1, so I will hopefully become more familiar with our digital cameras. Also, I will be attending Florida State's Film School in the fall, so SIGGRAPH 2000 should be an excellent way for me to put my foot in the door."

Kristina Bennett

"I became interested in SIGGRAPH when my teacher announced to my class that we had an opportunity to work with the best equipment available and some of the top minds in the business. I was able to attend meetings held at Franklin and at Kennedy where SIGGRAPH 2000 was discussed. I was also able to attend a special meeting for teachers at the New Orleans Museum of Art, where the speaker discussed how to bring computer graphics into the curriculum. Another student and I interviewed the speaker, Marty Sixkiller. He discussed how he became involved in the industry and how young people should get started. I'm excited about this summer. I can't wait to learn about all the new technology and techniques that are used today."

Alexis Brown

"I first joined the sigKIDS team near the beginning of the school year. I thought it would be an excellent opportunity to see what new ideas have been developed in the field of communications. My strengths in that area are currently digital and linear editing. I also have been studying film for two years under Mr. Werner. I logged the video for the NOMA project. I hope to learn more about SIGGRAPH in the next few months."

Luisito Caleon

"In a feeble attempt to further my future career in the oh-so-booming computer industry, I got a pass from Mr. Werner to attend the SIGGRAPH meeting to be held at sixth period. Call it concern for my future or my extreme hate for sixth period Spanish IV AP. A lot of people I did not expect to see were in attendance at the meeting. Contrary to my expectations, this seemed to be more than a reason to get out of class. The video we viewed was very well done in music-video form to some unknown techno song. It featured a lot of CG shorts I've seen here and there. Nikki arrived late to the meeting because of her Japanese class at the University of New Orleans but managed to get there before the meeting ended. We all were given Marc Barr's email address and the URL to the official SIGGRAPH 2000 site. I still was not quite sure how we were going to get involved in this massive convention. All I knew was that we would all get SIGGRAPH hats. Nothing like a good incentive to keep you going. Unfortunately, I had to return to class afterwards."

"November 1999: A sizeable group of Franklin students went to the SIGGRAPH outreach session at Kennedy High School. We were joined by a group of Kennedy "techies" who were skilled at computer hardware and basic networking. We also viewed a short on "The Physics of Cartoons." I, to this day, am still not quite sure what the purpose of attending the "outreach" was. My biggest worry was that I would have to go back to school early and take an English quiz I purposely did not prepare for, in hopes the "outreach" would take all afternoon. In lieu of returning to school, I went to Nikki's house to check my email. Upon arriving to school for play practice, I was informed everyone had congregated at PJ's Coffee & Tea in order to kill those extra minutes left in the school day."

"January 2000: The entire Ben Franklin sigKIDS community is currently running its base of communications through a mailing list on egroups.com, all thanks to Newhouse. Adam also created a message board and file storage systems on the site. Apparently, I have to watch what I say in the presence of the group (and with a certain Ryan Prechter)."

"February 200: Mr. Werner asked for volunteers to work on the first assignment for Ben Franklin sigKIDS. The assignment was to attend another outreach session with Marty Sixkiller of "ANTZ" fame at the New Orleans Museum of Art. I was unable to attend because of yet another play practice but agreed to help put the information from the outreach session online. After a few bouts of procrastination, I edited the text sent to me from Leslie Becker, and Nikki constructed yet another amazing Web site for the outreach session."

Gerrelde Ferrand

"I joined sigKIDS in January of 2000. I did not have the advantage of attending any of the seminars, but I heard that they were informative and innovative. The group that I am so pleased to be with created a wonderful Web site. The multimedia room at Benjamin Franklin has also allowed me to work with the following programs: Maya, Poser 4, True Space 4, and many other animation programs. I cannot wait to attend SIGGRAPH 2000. I am sure it will be an educational and interesting event that I will remember forever.

Adham Hussain

"When I first joined the sigKIDS team, I became introduced to 3D graphics. I started using a demo program called Studio Pro 2.5.3, and I did quite well with it. I made several 3D pictures that I also uploaded.

After I mastered the basics of Studio Pro, I decided to use some of the better 3D programs that were available to me. I first tried Bryce 3D, and I liked it, but I really wasn't interested in doing landscapes yet. I also tried using Maya, which is really great but very difficult to use. I attempted to make a bird fly. I had all of the polygons and animation set, but I just couldn't render it. Well, I didn't quite give up on Maya. I'll just come back to it later.

I have also done a bit of tampering with Poser 4. At first I didn't care for this program, but once I saw what can be done I became addicted. I now try to use this program as much as I can. So far, I have created several oddball animations, including a woman spinning up side down like a helicopter and a vampire with shifting emotions. My first goal is to make a character do a backflip kick and have someone getting hit by the kick. After I do this, I will try to create a one-minute fight scene, or I will try to create a character throwing a ball of fire from his hands (like in Street Fighter). This might be wishful thinking, but I don't care. I'll just keep on trying.

I first got involved with sigKIDS when I went to the presentation at Kennedy High School in November of 1999. I was instantly interested, because I always had a fascination with computer graphics, 3D animation, and video. I figured that this would be the perfect opportunity for me to express my interests, show off my talents, and use new programs that I was unable to use before. The presentation at Kennedy High School was very entertaining, and it increased my interest in 3D graphics."

Jeff Merlin

"I got interested in sigKIDS when Mr. Werner told our TV2 class that we were starting the program at school. I thought it would be good experience for me to work with the digital cameras and editing equipment, as well as the animation studios, and actually have a set goal in mind. I was also looking forward to attending SIGGRAPH 2000 in July and the opportunity to broadcast it over the Web, but unfortunately I will be out of town on those dates.

Thus far in sigKIDS, I have used Maya to program a ball bouncing through a fire hoop. I am also learning how to use Poser 4, and I have made a claymation video that will be digitized and uploaded to the Internet. I hope to accomplish much more in my last three weeks here at Franklin."

Adam Newhouse

- "• Beginning of the year: I play around with Bryce 3D and Ulead Cool 3D.
- Went to Spain over Thanksgiving vacation. I become interested in taking good pictures.
- I receive my digital camera for Christmas; I start taking more pictures, become photographer for Mu Alpha Theta.
- Learn the basics of cleaning up a picture using Photoshop; bad lighting at the awards ceremony for a Mu Alpha Theta competition allows me to hone my cleanup skills.
- Hash: Animation Master. After it came in, I played around with it some, created a few scenes using mostly preset objects and models I downloaded. Didn't really have time to learn it very well. Now that I have AM at home, I'm learning much more, as I have more time to go through the tons of tutorials that exist for it.
- TrueSpace 4.3. Began learning how to model things and create textures. Tried modeling the top of one of the desks at school as a beginner's project, but it never came out looking right. Created some realistic-looking floor tiles from school with my digital camera, Photoshop, and time spent getting the lighting right in a scene that had the tile texture repeating on a plane.
- Maya arrives. I open it up, poke around some, think to myself: "This is pretty cool." Start going through the first tutorial. Over the course of a few weeks, I learn many things, including the fact that you should never trust the screenshots in the Learning Maya 2 book. Currently trying to learn how to model, am about to skip the least three lessons on Tutorial 1 and go on to Tutorial 2, which looks like it involves a lot of modeling.
- I joined the sigKIDS team because I love computer graphics, and I've read a lot about SIGGRAPH, and thought it would be fun.
- NOMA presentation. This was an interesting project. It was neat seeing behind the scenes of "Antz" and hearing about it from someone who worked on the film. The interview afterward was nice too, because it allowed me to find out things you normally don't get to.
- Kennedy presentation. Fun. I liked the cartoon laws of physics."

Bianca Rey

"I'm most proficient in HTML, with computer graphics coming in second. Lately, I've been using Maya, Ulead Cool3D, and Poser. I've also been dabbling in Bryce 3D. I first got started with sigKIDS when I inquired about helping with the school Web page. So, I was referred to Mr. Werner through my Spanish teacher. I spoke to Mr. Werner about helping with the Web page, and he started telling me about sigKIDS, and the SIGGRAPH conference. That night I went on the SIGGRAPH Web site and became interested. In the next few weeks, I heard the announcement about the group looking for techies, and I went to the meeting. One of the seniors, Adam Newhouse, was in my study hall, and he brought me down to the Multimedia Room. Adam was the person who showed me around all of the programs. Eventually we began collaborating on different projects, although we could never agree on what. So I watched him for awhile, and then Maya came in, and we both knew the same amount. So we both began the first tutorial, which we called "bouncing the ball," and we started sharing information. So far I've only gotten the ball to jump through the hoop, but not really bouncing.

All the sigKIDS computer graphics people share a computer in the Multimedia Room, and they all seem to have my study hall. It's a race to use the computer, and the first one there gets it. Recently, I've picked up using Ulead Cool3d gif animator. I've been using all the different technologies to improve my Web site. I use HTML, DHTML, Javascript, and a little bit of CGI on my page. The most exciting thing about working with sigKIDS is all the technology I get to learn about, and all the animations and graphics the other members make. I've become friends with a lot of the graphics people, and I enjoy working with them. This summer, I hope to be working on the sigKIDS web site, and I can't wait to see what everyone has done with all the work they've been doing."

Jinger Roy

"I joined sigKIDS when my teacher, Mr. Werner, first told our TV class about it at the beginning of the school year. I thought that the combination of the Internet and the stuff I had learned in TV class would be right up my alley, considering the fact that I've been studying TV for two years now at Franklin. I've gone to most of the meetings this year, and I helped with one of the projects by submitting some of my photography for other people to use. I also contributed to the sigKIDS vault by putting up some of my HTML work for everyone to see, even though what I'm really interested in is the filming and editing aspect of SIGGRAPH. This summer, I hope to help with volunteering for SIGGRAPH 2000 as much as possible, probably mainly in the areas of filming and editing with digital video."

Nikki Slaughter

"sigKIDS... what can I say? It's an unbelievable opportunity. I've been interested in graphic design since my freshmen year and plan to make it my major in college. When I heard about the sigKIDS project, I quickly sprang to it. Just being involved in the sigKIDS project has motivated me to learn more about programs like Adobe Photoshop and Dreamweaver. I also can't wait to attend SIGGRAPH 2000 this summer.

My last completed sigKIDS assignment was making the Web page for the NOMA presentation. The programs I used were mainly Photoshop 5.5, ImageReady, and Dreamweaver 3. All great programs and... not too hard to get the hang of. It was actually my first time using Dreamweaver, and even though I never cracked open a user's manual, I think I did a pretty good job!

My favorite part of that project: I had a lot of fun recreating SIGGRAPH 2000's header image for one thing! If you don't stare at it too hard you'd barely notice it was a replica. Second, I loved making those rollover images. I don't know where I got the idea of making them like TVs, though. Maybe I had been watching the TV too much that night? And speaking of rollovers, my third favorite part was experimenting with Dreamweaver. It was so easy to make rollovers! I created the Web pages before the actual templates I was supposed to use were uploaded. So without the template I was clueless as to how I was going to make the images I downloaded from sigKIDS site work as rollovers. I'm happy I didn't have the templates though. It was a fun challenge to recreate the effects (even if I didn't type the HTML code on my own). I think I got pretty close to the original sigKIDS templates.

Hmm... new topic: I've been using FrontPage for my personal site, otakuonline.com, since the FP '98 version. Before that, I did my own coding. So why did I take the lazy way out and rely on a program for the coding? Good question! Well, over time I realized that although doing my own coding was fun, I enjoyed working on the graphics and layout much more. Using a program like FP allowed me to focus more on the look of my site. Unfortunately, FrontPage's HTML coding can get pretty messy, especially with all its special "web bots" and whatnot. FrontPage also tends to make the site look messy on Netscape browsers. Recently, I started playing around with Dreamweaver 3. What a find! It gives me more control of the actual HTML but still allows me to put most of my focus on the design. The best thing is that its codes are clean and it's quite easy to make your site compatible for most browsers. I've also tried Adobe's GoLive. It has a lot of nice features. I haven't really used it much, so I can't really give a solid opinion on it.

Right now, I am learning the basics of Flash 4. After I get a little more time with Flash, I'm going to work a little with Fireworks 3. As you can see, I'm obsessed with Macromedia and Adobe software. I hope our sigKIDS program will offer some good instructional meetings for some of their programs."



*Benjamin Franklin High School sigKIDS team
Adam Newhouse, Nikki Slaughter, Luisito Caleon, Bianca Rey, Matt Smith,
Ryan Prechter, Alexis Brown, Adham Hussain, Paul Werner, Leslie Becker*

Two model boats are remote-controlled using the arrow keys of a computer keyboard or a joystick. Users engage in a slow-motion soccer encounter on the water, pushing the ball with their boats. Some very

basic initial rules of the game are set, but the level of difficulty can be increased as users become more proficient in anticipating the movements of the boat in response to their instructions.



RVV (Remote Visual Vehicle)

Gevel Marrero
William Sayer
Pratt Institute

Luis Marrero
Northeastern University

Equipment Sponsor
Pratt Institute

The RVV is an interactive remote vehicle that provides realtime video feedback to the operator. It can be used to explore dangerous environments or just for fun. This project is an exploration of existing technologies and what can be done when they are taken to a different level.

The operator controls this hands-free, motion-controlled vehicle via video feedback through stereoscopic glasses and a head-tracking system mounted on the RVV and connected to a desktop computer. The vehicle itself is an out-of-the-box product with a micro camera and a transmitter that sends the wireless video images to the operator.





Contents

- 173 3D Plenoptic Functions with Cylindrical Maps
- 174 4D Swathing to Automatically Inject Character into Animations
- 175 Acting Rehearsal in Virtual Reality
- 176 Adaptive Visualization of Large Dynamic Non-Manifold Models
- 177 Artistic Renderings of Dynamic Motion
- 178 Autostereoscopic Display with Real-Image Screen
- 179 Behavioral Animation by Constraint-Based Adaptive Search
- 180 Better Blending of Implicit Objects at Different Scales
- 181 A Branching Cellular Texture Basis Function
- 182 Bresenham Noise
- 183 Bringing the Rock Monster to Life for “Galaxy Quest”
- 184 Building A Digital Actor From The Inside, Out
- 185 Building “Shrek’s” Natural Environment (Trees, Plants, and Grass)
- 186 Building Usable Objects for Virtual Environments
- 187 Bunshin: A Believable Avatar Surrogate for Both Scripted and On-the-Fly Pen-Based Control in a Presentation Environment
- 188 Cephalic Visitors
- 189 “City of Brass:” Hieroglyphic Perceptions in Digital Cinema
- 190 Clothing Simulations in “Shrek”
- 191 Combination of Two- and Three-Dimensional Space for Solid Modeling
- 192 Compressing Texture Maps for Large Real Environments
- 193 Conception To Birth
- 194 A Constructive Approach to Playing Dynamic Images
- 195 Corrective Texturing
- 196 Directing The Digital Actor: Real World And Digital Challenges
- 197 Displacement Mapping Rendering Hardware using Adaptive Tessellation
- 198 Dynamic Animation and Control Environment (DANCE)
- 199 Efficient Out-of-Core Build of a Topological Data Structure
- 200 Enhancement of Virtual Interaction Devices Using Pseudo Force Sensation
- 201 Evaluating Image Quality Metrics vs. Human Evaluation
- 202 Fast Multi-Resolution Operations on Images in the Wavelet Domain
- 203 Fitting Catmull-Clark Surfaces to Laser Scanned Data from 3D Models
- 204 Flexible Application Design Using VR Juggler
- 205 From Motion Capture to Motion Texture
- 206 GeoSCAPE: 3D Visualization of On-Site Archaeological Excavation Using a Vectorizing Tape Measure
- 207 Gesture-Based Ping-Pong Game Using Real-Time Depth-Image Input Device
- 208 Getting “Shrek” Muddy
- 209 GINGA
- 210 Going One on One in “Raptor vs. Raptor”
- 211 Hardware-Accelerated Silhouette Matching
- 212 Heuristic Backface Culling of Animated Subdivision Surfaces
- 213 i2TV: Experiments in New Formats for Connected Communities
- 214 Illustrating Uncertainty
- 215 Implementing the OpenGL API with Direct Rendering in Linux
- 216 An Insight into Physiquing a Digital Human for “Hollow Man”
- 217 Integrated Pose-driven Character Dynamics
- 218 Interactive 3D Rendering on PDA Platforms
- 219 Interactive Rendering of Shafts of Light Using a Hardware-Accelerated Volume Rendering Technique
- 220 “la_alma” Interactive Book
- 221 Leveraging the Team Across the World: Creature Development for “Pitch Black”
- 222 Line Art Rendering Using a Modified A-Buffer
- 223 Liquid Time: An Exploration of Video Cubism
- 224 A Low-Cost, Portable Light Field Capture Device
- 225 MAJOR DAMAGE: The Virtual Garage Studio Processes for Producing Animations
- 226 A Match Moving Technique for Human Body Images
- 227 “Meshsweeper:” Fast Closest Point on a Polygonal Mesh
- 228 A Method Of Generating Pavement Textures Using The Square Packing Technique
- 229 The Met
- 230 The Metric Cursor
- 231 The Mixed Reality Stage
- 232 Modeling with Guts: The Complete Replication of an Actor’s Biology
- 233 MorphAd: A New Concept in Product Advertisement for E-commerce
- 234 MorphZhapas
- 235 Motion Adaptation with Self-Intersection Avoidance
- 236 The Museum as Information Architecture

- 237 Natural Interfaces via Real-Time Video
- 238 The NAVE: Design and Implementation of a Non-Expensive Immersive Virtual Environment
- 239 Near Real-Time Cutting
- 240 On Determining the Haptic Smoothness of Force-Shaded Surfaces
- 241 An Optical See-Through Display for Mixed Reality
- 242 Painting Space with BLUI
- 243 Panoscope 360^o
- 244 Polytope Sections
- 245 Presence and Task Performance in Virtual Environments: A Reality Centred Approach
- 246 The Puppeteer Behind The Avatar
- 247 A Realistic 3D Spatial/Symbolic Model of the Visible Human
- 248 Real-Time Human Hair Modeling and Animation
- 249 Real-Time Interactive Grid-Free Fluid Dynamics
- 250 Real-Time Painterly Rendering
- 251 Reconstruction Of The World's First Panoramic Immersive Virtual Environment Using Contemporary VR Technology
- 252 Robust Invisible Watermarking of Volume Data
- 253 The Role of Eye Gaze in Avatar-Mediated Conversational Interfaces
- 254 Scalable Motion Simulation
- 255 Shedding Light on Image-Based Rendering
- 256 Six layers of Kevin Bacon: A Digital Human Going Invisible
- 257 Sleepy Hollow: Digital Blood (Dynamic) Simulations and Volumetric Rendering
- 258 Snap To It! Automatic 3D Object And Silhouette Registration
- 259 A Solution to Vertex Path Problem in Shape Blending by Using a Dependency Graph
- 260 Sort-First Parallel Rendering with a Cluster of PCs
- 261 Staging the Space of Mixed Reality
- 262 "Superbodies"
- 263 Survivors: Testimonies of the Holocaust: Humanity through Technology
- 264 Tailored Virtual Tours In Cultural Heritage Worlds
- 265 What is "Tal vez Mañana"?
- 266 Texture Mixing and Texture Movie Synthesis using Statistical Learning
- 267 Time Document: New Media for Dynamic Public Art
- 268 TimingChart: Timing Animation via Traditional Methods
- 269 "Toy Story 2" Technical Challenges
- 270 The Use of Spryticle in the Visual FX for "The Road to El Dorado"
- 271 The Use of Traditional and 2D CG Effects as Used to Create the Magic in "The Road to El Dorado"
- 272 Virtual Human That Can Predict Comfort and Pain Level
- 273 A Virtual Reality Interface for Analyzing Remotely Sensed Forestry Data
- 274 Visapult: A Prototype Remote and Distributed Visualization Application and Framework
- 275 Visibility Driven Rasterization
- 276 Volumetric Modeling of Additional Techniques in Colored Pencil Drawing
- 277 VScape: Autonomous Intelligent Behavior in Virtual Worlds using UML
- 278 Water Dome Project
- 279 An XML Render e-Service

SKETCHES & APPLICATIONS

TOM APPOLLONI
HARRIS CORPORATION



Committee

CHAIR

Tom Appolloni
Harris Corporation

COMMITTEE

Paul Debevec
USC Institute for Creative Technologies

Jacquelyn Martino
Philips Research, USA

Scott Senften
SGI

JURY

Ronen Barzel
Pixar Animation Studios

Christoph Bregler
Stanford University

Anna Maria Chupa
Mississippi State University

Alan Commike
SGI

Becky Fuson

Hugues Hoppe
Microsoft Research

Michelle Meeker
Pacific Data Images

Jacquelyn Ford Morie
USC Institute for Creative Technologies

Ken Musgrave
FractalWorlds

Doug Roble
Digital Domain

Brian Rosen
Pixar Animation Studios

Chris Shaw
Georgia Institute of Technology

Dena Slothower
Stanford University

Annette Weintraub
City College of New York

Brian Wyvill
University of Calgary



THE CHALLENGES ARE TO EXPAND OUR HORIZONS, to make us look at problems from new angles, and to make us rethink and possibly redesign the way we work. The forum for presenting these new concepts is Sketches & Applications. Here, we present works in progress, initial research results, new approaches to solving old problems, and new processes to advance how today's computer graphics and interactive techniques are moving from theory to everyday use.

Sketches & Applications is the program where all the disciplines associated with computer graphics and interactive techniques are intermingled. The presentations selected this year are not separated by discipline but combined by common themes. The goal of the Sketches & Applications committee was to merge art and science, design and animation, application and multimedia into sessions that intrigue as well as inform; challenge as well as enlighten. These sessions are designed to bring together members of all the disciplines in order to spark discussion and broaden horizons.

The jury received almost 300 submissions. After many hours of review, discussion, and deliberation, it selected 107 presentations in three central categories:

ANIMATION

Animation Sketches & Applications explain how technology and artistry have combined to create some of this year's most compelling effects, stories, characters, and visualizations. Whether the sketch is presented by a major studio or independent animators, there is a great deal to learn about techniques, aesthetics, and project coordination. This year's selections include the latest effects and technology work for blockbuster films as well as amazing results produced by small animation houses and independents.

ART, DESIGN, AND MULTIMEDIA

These Sketches and Applications demonstrate how artists, designers, and engineers are pushing computer graphics as they create new tools and forms to express themselves as well as apply known techniques in elegant and innovative ways. This year's jury selections range in topic from space as a carrier for information to explorations in public art, complemented by insights into narrative forms and approaches to 3D graphics.

TECHNICAL

Technical Sketches and Applications continues the heritage of sketches by presenting the entire gamut of computer graphics and interactive techniques from machine-level implementations to VR applications, from initial research findings to the lessons learned from productions. We're rendering with hardware, textures, and the Web. Avatars are moving autonomously while humans use new 3D techniques for interaction. We look at the approaches for both scientific and artistic visualization. We examine motion in the form of capture, haptics, and dynamics.

The Sketches and Applications Committee would like to thank everyone who submitted their work for consideration. Your continued support with new ideas and new insights from throughout the international community is an affirmation of the continued drive and enthusiasm for the advancement of computer graphics and interactive techniques.

I would like to thank the members of the Sketches & Applications Committee as well as the SIGGRAPH 2000 Conference Committee for maintaining such an encouraging and collaborative environment. My personal gratitude goes out to Vicki Caulfield for keeping all of us on the right track as my program administrator and to Carrie Ewert for coordinating the publication process. Without them, we could not have reached what we set out to achieve. Lastly, I would like to take this opportunity to thank Jackie White and Jennifer Anderson for their encouragement for all the SIGGRAPH 2000 program chairs to create their visions, my colleagues at the Harris Corporation for the support to create Sketches & Applications, and, most importantly, my wife, Candy, for putting up with my obsessions.

A plenoptic function is a 5D function that defines a radiance value at any point (V_x, V_y, V_z) in the space and in any direction (θ, ϕ) . This sketch presents a novel 3D representation of a plenoptic function that uses cylindrical maps placed on a circle. This representation resolves the drawbacks of concentric mosaics,¹ a 3D representation of a plenoptic function proposed recently.

PROPOSED REPRESENTATION OF A PLENOPTIC FUNCTION

The 3D representation of a plenoptic function proposed in this sketch consists of cylindrical maps placed on a circle (see Figure 1). As in the lumigraph, light-field rendering, and concentric mosaics, we assume that the value of a plenoptic function is constant on a line. The three parameters of the representation are the rotation angle θ , direction ϕ , and elevation l . The radius R of the circle is fixed. Figures 2(a) and 2(b) show the simplified versions of the representation that handle looking outward from an environment and looking inward at an object, respectively.

COMPARISON WITH CONCENTRIC MOSAICS

The representation proposed in this sketch has benefits over concentric mosaics in three ways. First, the construction of an environment map

is much easier than that of a mosaic. An environment map can be constructed from a sequence of images taken by a non-special camera. In contrast, we need special hardware such as a slit camera to obtain a mosaic. Second, concentric mosaics are designed for looking outward from an environment, and it looks almost impossible to apply them to an object looking inward. However, the proposed representation naturally supports the looking inward, as well as an environment looking outward. For example, we can look into an object from any point outside of the circle in Figure 2(b). Third, the depth distortion can be corrected more easily with the proposed representation than with concentric mosaics.

EXPERIMENTAL RESULT

Figure 3(a) shows a cylindrical map used to construct the representation of a plenoptic function. The images in Figure 3(b) are two examples of looking inward generated from the representation.

www.postech.ac.kr/~jin
www.postech.ac.kr/~leesy

Reference

1. Shum, H.Y. & He, L.W. (1999). Rendering with concentric mosaics. Computer Graphics (Proc. SIGGRAPH 99), 299-306.

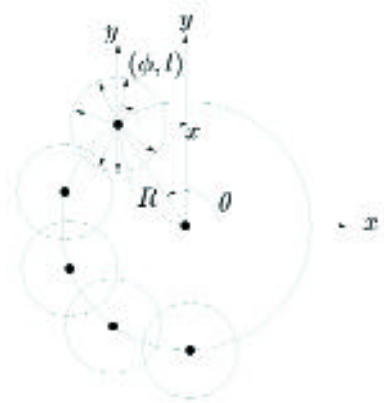


Figure 1: Proposed representation of a plenoptic function.




Figure 2: Simplified versions of the representation. (a) looking outward, (b) looking inward.



Figure 3: Input cylindrical map and rendered images.

Animation packages are good at making physically realistic animations. However, they lack the ability to automatically inject cartoon-style effects such as extreme deformations of limbs. In this sketch, we present a technique for automatically introducing deformations into the animation of rigid objects, adding character.

TRANSFORMATION

We inject flexibility by applying a transformation over the scene prior to rendering. We transform the scene by keeping the bottom (for example) of the scene unchanged, while rendering the top of the scene slightly ahead in time, with a continuum in between. In this way, a rigid object that is translated and rotated over time in the scene will bend slightly due to its non-linear motion across the time field. We call this process time swathing.

There is an analogy here with the use of spatio-temporal cubes in computer vision. This process consists of forming a 3D cube of data by stacking 2D still images (each from a different moment in time), whereupon a 2D slice can be taken from this cube. Our approach is similar except that we define a 4D scene (consisting of the 3D objects animated in time), from which we take a 3D slice that is then rendered using an ordinary projection.

IMPLEMENTATION

The most efficient way to implement time swathing is by defining geometries in terms of a 4-dimensional space, representing the X, Y, Z, and time coordinates. The time swath is taken by slicing a 3D volume of the 4D space and then rendering the 3D volume. Prior to slicing, the 3D volume is rotated and stretched relative to the time axis. For every frame, we take a different swath from the 4D space, typically by moving the 3D volume along the time-axis. If the 3D volume is perpendicular to the time axis, a normal rendering results. When rotations around

the T-X, T-Y and T-Z pairs of axes are defined, then different swath types, and hence different effects, result.

We convert traditional 3D NURBS animated over time into 4D NURBS. 4D NURBS work just like their 2D and 3D counterparts, except that they are described in terms of (u,v,w) spanning points. Using extended NURBS mathematics, slicing a 3D volume through 4D NURBS results in traditional 3D NURBS, with 2D sets of spanning points. We subsequently render these 3D NURBS using a conventional ray tracer.

EFFECTS

The subtle swathing of 3D objects adds “character” to those objects. Rigid geometries tend to bend smoothly, as shown in Figure 1: a cartoon character on a stepper-exercising machine. The swathed stick-man is more cartoon-like and is produced quickly and automatically.


More advanced forms of swathing take advantage of the fact that swathing does not have to be applied to the whole scene and does not have to be linear in time. For instance, in Figure 2 a swath rotates around the X-axis, while the cog spins around the Z-axis. This leads to a flexing of the plane of the cog and non-linear deformation of the cog’s teeth.

Swathing can be bound to a particular part of the scene. For instance, in the stick-man example, the steps could be left unswathed.

www.cs.bris.ac.uk/

ACKNOWLEDGEMENTS

We are grateful for the help of Dave Blackham and Dave Gibson of Broadcast Film and Video, Bristol.



Stick-Man before (left) and after (right) a linear swath. Note the bending of the steps as well as the suppleness of the limbs. The legs change length (this frame) and bend (subsequent frames). In the animation (www.cs.bris.ac.uk/Research/Digitalmedia/Swath/), one can see the stretching and bending of stick-man’s spine in addition to the deformations of the limbs.



The Cog before (top left) and after (bottom right) swathing. Note the non-linear deformation of the teeth. In the animation (URLs on the left), the teeth bend fluidly in the direction of the cog’s rotation, while the entire cog is flexed in its plane due to the rotation of the swath itself.

Three pairs of professional actors and a director met in a shared non-immersive virtual reality system over a two-week period for about four and a half hours to rehearse a short play. The actors and director did not meet each other physically until a short time before they presented a live rehearsal to an audience. The virtual rehearsal system was implemented using the DIVE 3.3x platform (www.sics.se/dive), which has built-in networking support including spatialized voice communication. The study examined the extent to which virtual reality could be used by the actors and director to rehearse their later live performance. Four indicators elicited by questionnaires show that over the rehearsal period their sense of presence in the virtual rehearsal space, their co-presence with other actors, and their degree of cooperation all increased. More importantly, their evaluation of the extent to which the virtual rehearsal was similar to a real rehearsal also increased.

Both actors and the director used full-screen desktop displays. The actors were embodied as avatars who could see and talk to one another in the shared virtual space. The director was a disembodied voice with no visual representation. The actors controlled their facial expressions with mouse strokes on a “smiley” face. To evoke a smile, they would draw one, and they could achieve other facial expressions by stroking a mouth and eyebrow shape combination. They could use the mouse turn their heads. The actors fully exploited this interactive interface. For example, they learned without prompting how to make “envelope” gestures such as small nods and head shakes with the mouse interface.

From free-format discussions held with each of the actors separately after the sessions, it was clear that the actors were initially doubtful about the utility of the system for rehearsal. Subsequently, they became very much at ease with the system and had positive comments about the virtual rehearsal process. Figure 1 illustrates how the virtual space was used in accordance with the desired emotional content. Figure 2 illustrates the mapping of spatial arrangements from the virtual rehearsal to the real rehearsal.

An independent observer (David Chambers of the Yale School of Drama) noted: “It was remarkable watching the actual actors on separate monitors as they rehearsed on-screen ... When present, the emotional content was clearly perceivable, vocally and gesturally.” There remains a need to compare the results with traditional rehearsal and video-conferencing.

ACKNOWLEDGEMENTS

We would like to thank Ivor Benjamin for his help during the rehearsal experiments. We would also like to thank Mitch Mitchel for his post-production work on the documentary video. This project is funded by the UK Digital VCE.



Figure 1: The actors varied spatial distances between each another according to emotional needs




Figure 2: Blocking was achieved: the same moment in a virtual and real performance.

Many disciplines require large high-resolution geometric models, which exceed hardware capacities. In particular, geoscientists want to build models of the earth's crust for oil field reservoir characterizations. Models are usually non-manifold and are built interactively and incrementally using Boolean operations, just as in CAD, from surfaces generated from acquisition data. The models are dynamic, because Boolean operations modify the surfaces in the model. Multiresolution techniques^{1,2} are used to represent large static models, such as pre-built CAD models or meshes generated from a 3D scanner. With existing approaches, when the model is dynamic, the multiresolution representation has to be reconstructed from scratch. To tackle these problems, we have developed SIGMA (scalable interactive geometric modeling architecture).^{3,4} This sketch presents our new results, which include building and decimating more complex models and adaptive visualization of physical material properties together with the model.

The essence of SIGMA is a multiresolution surface representation implemented as a quadtree. The representation supports both Boolean operations and adaptive visualization. Figure 1 shows the data flow using SIGMA. During the modeling process, for example, when a new surface is inserted into the model via Boolean operations, SIGMA incrementally updates, instead of reconstructing, the surface representations.³ To support visualization, the model is decimated in a view-dependent manner. SIGMA traverses quad trees of the surfaces in the scene to determine vertices that will contribute to the decimated

model. Intersections and boundaries of the surfaces are represented by a collection of edges. Constrained Delaunay Triangulation is used to triangulate selected vertices and edges in the parameter space of a surface. Edges of intersection curves are shared by surfaces during triangulation to prevent cracks in the decimated model. Figure 2 illustrates how we apply adaptive multiresolution visualization during the modeling process. Figure 3 shows some snapshots that illustrate the modeling process and the dynamics of the model. The model starts with a volume of interest. Surfaces are incrementally inserted into the model. Each insertion may split volumes into subvolumes. We apply adaptive visualization to the model during this process to enable interactive modeling. The visualization has surface mode and volume mode, which renders the surfaces and resulting volumes respectively. A user can assign physical material properties to surfaces and volumes during or after the model is being built. We render these properties by mapping property values to colors.

References

1. Hoppe, H. (1996). Progressive meshes, *Proc. SIGGRAPH 96*, 99-108.
2. Lounsbery, M., DeRose, T.D., & Warren, J. (1997). Multiresolution analysis for surfaces of arbitrary topological type. *ACM Transactions on Graphics*, 16 (1), 34-73.
3. Hammersley, R., Lu, H.Q., & Assa, S. (1999, August). Geometric modeling with a multiresolution representation. Eleventh Canadian Conference on Computational Geometry.
4. Lu, H.Q. & Hammersley, R. (2000, February). Adaptive visualization for interactive geometric modeling in geoscience, WSCG'2000.



Figure 1: Data flow in SIGMA approach.




Figure 2: The role of adaptive visualization.

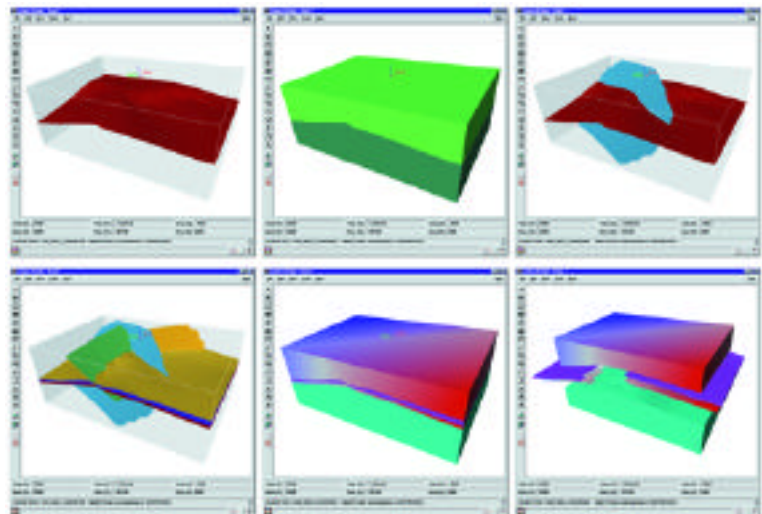


Figure 3: Snapshots from a modeling process (left to right, top to bottom).

Traditional arts such as painting, sculpture, and photography have long emphasized the dynamics of human motion. To answer the question of whether computer-rendering techniques can similarly enhance our perception of the dynamics of motion, we developed six rendering styles and performed a user study on one of the styles to determine if it affected the viewer's perception of an animated character's velocity.

We adapted six traditional techniques from hand animation, painting, sculpture, and photography into rendering styles and used them with motion from a dynamically simulated runner. For example, multiple exposures on a still photograph depict the different positions of an object over time via stills. We developed a "shadow" technique that simultaneously showed the animated runner's current position and a set of previous positions (figure e). A "geometric" rendering algorithm was developed based on the style of Duchamp's "Nude Descending a Staircase." The physical aspect of the model was abstracted into many overlapping geometric shapes (figure f). Similarly, our rendering styles of motion lines, motion blur, distortion, and unbalanced figures were motivated by cartoon lines, camera shutter speed, and the exaggerated muscles seen in Futurism (figure c).

Because velocity is an easily quantifiable measure of dynamics, our user study assessed whether the viewer's ability to differentiate between the speed of simulated running motions was affected by the rendering style. Subjects viewed two motion sequences rendered with either the same or different styles and made a judgment as to which was faster.

The motion sequences were generated from six trials of a simulation of a human runner with speeds of 4.5, 4.7, 4.8, 4.9, 5.0, and 5.1 m/s. We tested one rendering style, the multiple exposure or "shadow" technique against the "normal" runner rendered with a standard polygonal model. The test contained 54 pairs of movies rendered from the motion sequences. These movies were separated into three categories with 18 pairs in each: 1) two "normal" runners, 2) a slower "normal" runner and a faster "shadow" runner, and 3) a faster "normal" runner and a slower "shadow" runner. The of the 54 pairs were shown in random order.

Because only the speed and the rendering style changed, we were able to test for correlation between rendering style and perception of speed. For each pair of motions, viewers were asked to make an A/B comparison and decide which runner was moving faster. We conducted paired t-tests to determine if our results were statistically significant. A paired t-test reports the probability that the actual mean difference is consistent with zero and tells us if a subject's perception of speed changed between the two categories.

The number of questions incorrectly answered in category 1 was used as a baseline in the paired t-tests. Paired t-tests of category 1 and the combination of categories 2 and 3 determine if the viewers were better able to judge the speed of the runner when rendered normally or with the "shadow" technique. This test indicates a 90 confidence that user perception was changed. The paired t-test for categories 1 and 3 determines if viewers were inclined to perceive the shadow runner as moving more quickly, and the paired t-test for categories 1 and 2 determines the opposite. The paired t-test of categories 1 and 2 was statistically insignificant, indicating that the shadow technique did not decrease

perception of velocity. With categories 1 and 3, there was a statistical significance at 95 confidence level, indicating that viewers tended to perceive the shadow runner as moving more rapidly than was the case. This result indicates that the shadow rendering style makes the runner appear to be moving more quickly.

Results of a paired t-test on category 1-(2+3) for 17 subjects t = 1.83 Assuming null hypothesis, probability of this result = 0.08 Statistically significant at 90% confidence level.
Results of a paired t-test on category 1-3 for 17 subjects t = 3.57 Assuming null hypothesis, probability of this result = 0.003 Statistically significant at 95% confidence level
Results of a paired t test on category 1-2 for 17 subjects t = 1.57 Assuming null hypothesis, probability of this result = 0.136 Statistically significant at 90% confidence level

Although we have rigorously tested only the shadow technique, we intend to test the remaining rendering styles. Based on preliminary results, we expect that the "enlarge" technique will make the runner appear to be moving more rapidly, while the "geometric" technique will make the runner appear to be moving more slowly. Other styles might also affect the viewer's perception of the dynamics and could be tested in a similar manner.

- Hodgins, J.K., Wooten, W.L., Brogan, D., & O'Brien, J.F. (1995). Animating human athletics. *Proceedings of SIGGRAPH 95*, 71-78.
- Hodgins, J.K. (1998). Simulation of human running. *IEEE International Conference on Robotics and Automation*, 320-325.
- Hodgins, J.K. (1998). Perception of human motion with different geometric models. *IEEE International Conference on Robotics and Automation*.
- Johansson, G. (1973). Visual perception of biological motion and a model for its analysis. *Perception and Psychophysics*, 14, 201-211.
- Cutting, J.E. & Kozlowski, L.T. (1977). Recognizing friends by their walk: Gait perception without familiarity cues. *Bulletin of the Psychonomic Society*, 9, 353-356.
- Tisdall, C. & Bozzolla, A. (1978). *Futurism*. New York: Oxford University Press.
- Dagognet, F. & Marey, E. (1992). *A passion for the trace*. Cambridge: MIT Press.
- Rubin, W. (Ed.). (1980). *Pablo Picasso: A retrospective*. New York: The Museum of Modern Art.
- Janson, A.F. (1991). *History of Art*. Englewood Cliffs, New Jersey: Prentice Hall.




Figure 1: a) line, b) blur, c) enlarge, d) none, e) shadow, and f) geometric runner.

A new type of autostereoscopic 3D display is proposed. It consists of Fresnel lenses, a conventional display that presents two polarized images, polarizing filters mounted on a mobile stage, a head-tracking device, and a computer that controls these devices. The feature of the proposed system is that it generates a real image in front of the viewer, which reduces contradiction between the focus information and the parallax information when a 3D image is to be presented near the viewer. The detail of the proposed system is shown in Figure 1. Here, Fresnel lenses present the real image of the display in front of the viewer. Also, because of the lenses, the light inputs to the right eye and the left eye gather at different points between the display and the lenses. There, the lights can be filtered so that different images can be observed in each eye. Since the point where the light input gathers shifts as the observer moves, the position of the filter is transferred to the proper point as the observer moves. Thus the viewer can see images suitable for both eyes without wearing stereoscopic glasses.

To make suitable images for undistorted 3D perception, it is necessary to take into account the optical path, including the refraction by the Fresnel lenses. For example to present triangle ABC in Figure 2, triangle A'B'C' for the right eye and triangle A''B''C'' for the left eye must be presented on surface S. Here vertices A', B', C', A'', B'', C'' are on the points where the optical paths including line segments E'A', E'B', E'C', E''A'', E''B'', E''C'' intersect with surface S. This mapping of 3D space on surface S depends on the position of the eyes E' and E'' and has to be updated when the observer moves. To realize real-time mapping updates, the mapping function must be executed fast enough. We have succeeded in arranging the display and the lenses so that the mapping function can be approximated to be linear, which enables real-time display of almost undistorted 3D space. The proposed system can be used as an autostereoscopic 3D workbench where viewers can interact with 3D images directly with their hands without suffering from the severe eyestrain caused by the great contradiction between focus and parallax. Therefore, this workbench is applicable to VR training systems or human interfaces for teleoperation.




Figure 1: System configuration of the proposed autostereoscopic display.




Figure 2: Image generation to present undistorted 3D space.

Our aim is to design a high-level language for describing behaviors for autonomous agents in virtual worlds, together with an efficient run-time algorithm to implement those behaviors. The basis of this declarative language is the notion of constraint,⁵ which can be used to represent goals that the agents are trying to achieve. Behaviors are expressed as sets of goal constraints that have to be satisfied by the agent. The agent thus maintains a set of goal constraints, and we define a simple but effective framework, extended and abstracted from Codognet³ based on the formalism of Constraint Satisfaction Problems (CSP).

As a first application of this framework, we have considered the problem of navigation of autonomous creatures in virtual worlds. More generally, our framework can be used as a motivation architecture for such virtual creatures, by considering variables for denoting internal states (for example, energy, thirst, etc) and goal constraints for defining internal needs (for example, the energy should stay above a certain level), routine behaviors (if the energy falls below some level, go for food), or external desired properties (for example, stay away from predators).

The basic spatial constraints for autonomous navigation are:


<i>Constraint</i>	<i>Declarative meaning</i>
In(Region)	Stay within the zone defined by Region
out(Region)	Stay outside the zone defined by Region
go(Object)	Move toward the location of Object
away(Object)	Move away from the location of Object
Attraction(Stimulus)	Move toward the source of the stimulus
Repulsion(Stimulus)	Move away from the source of the stimulus

It is clear that a combination of such goal constraints could produce quite complex behaviors (for example, that the agent should go toward some object, avoid all objects it perceives, and stay away from some predefined regions). A following behavior can be simply obtained as a combination of a go constraint (to move toward the followed agent) and an out constraint (to stay at a certain distance).

In order to solve a conjunction of constraint goals, we have designed an adaptive-search mechanism based on local search techniques.¹ Indeed, one can define, for each goal constraint, a predefined repair mechanism that will propose (in case the constraint is not satisfied) an action that could reduce the degree of violation of the constraint, based on an “error” function that will indicate how much the constraint is violated. To treat conjunction of goals, Adaptive Search relies on iterative repair based on constraint errors information combined in variables (the error associated with a variable is the sum of the errors on constraints in which it appears). We use a heuristic method in seeking to reduce the error on the worst variable.

References

1. Aarts, E. & Lenstra, J. (Eds). (1997). *Local search in combinatorial optimization*. New York: John Wiley and Sons.
2. Badler, N., Philips, C., & Webber, B. (1993). *Simulating humans: Computer graphics animation and control*. Oxford University Press.
3. Codognet, P. (1999). Declarative behaviors for virtual creatures. *Proc. ICAT'99*, 8th International Conference on Augmented Reality and Tele-existence.
4. Koga, Y., Becker, C., Svihura, M., & Zhu, D. (1998). On intelligent digital actors. *Proc. Imagina 98*.
5. Saraswat, V., Van Hentenryck, P., et al. (1996). Constraint programming. *ACM Computing Surveys*, 28 (4).



Skeletal implicit surfaces¹ are useful for building models, particularly where smooth blends are required. Such models are generated as iso-surfaces in a scalar field. Blending called “proximity blending” is achieved by placing field generators (skeletal primitives) in close proximity so that their fields overlap.

The top row of Figure 1 indicates an interesting problem with proximity blending. The figure shows a series of close-ups of a point primitive A, with a large radius blended with a point primitive B, with a much smaller radius. It can be seen that the shape of the blend region is approximately equivalent to using union between the primitives and does not have the smooth blending shown in the close-up views in the bottom row of Figure 1. The effect is increased as the ratio of radii increases.

Figure 2 shows the iso-value contours of B, A+B, and the zero surface of B. In this case, the ratio of the radii of the large to the small sphere is 8:1. The left hand of Figure 2 is the regular blending method, and it can be seen that the iso-surface of A+B is very close to the zero-surface of B. This is due to the functions used to calculate the field values for A and B, which are shown in Figure 3. It can be seen that the values returned for A overwhelm the contributions from B.

To improve the shape of the blend, we propose a method we call “restricted blending.” We modify the distance over which the field function² drops from one to zero (the “blending distance”) for the larger primitive, to match that of the smaller. The field function of A is scaled,

then translated, resulting in the field function shown as the blue line in Figure 3. The field function is finally clipped so that $f(r)=1$ if $r < (R_A - R_B)/2$ and $f(r)=0$ if $r > (R_A + R_B)/2$.

The right hand of Figure 2 shows a dramatic change in the position of the contour representing the sum of the two primitives when using restricted blending. The resulting blends for various ratios of radii are shown in Figure 1. The field function has been restricted to vary within the blending distance of the smaller sphere in each case. The colour at any point on the surface is calculated by weighting the colours of the contributing primitives by its field value. Restricted blending also allows the colour of the smaller primitive to affect the result as can be seen in Figure 1.

With primitives other than the sphere the same approach applies. Restricted blending works well for two primitives by choosing the appropriate blending distance. Unfortunately this approach does not necessarily work when it is necessary to blend more than two primitives. We are experimenting with continuous functions that adjust the effective blending distance of each primitive in the vicinity of other.

See also: www.cpsc.ucalgary.ca/~blob/papers2000.html

1. Bloomenthal, J. (Ed.). (1997). *Introduction to implicit surfaces*. Morgan Kaufmann, ISBN 1-55860-233-X.
2. Wyvill, G., McPheeters, C., & Wyvill, B. (1986). Data structure for soft objects. *The Visual Computer*, 2 (4)




Figure 1: Point primitives of decreasing radii, regular blend (top), using restricted blend (bottom).




Figure 2: Zero surface of B dotted black line is iso-surface of B alone.




Figure 3: Field functions.

A cell-basis function can create intricate textures that simulate many patterns seen in nature. Worley⁵ shows examples of complex and interesting patterns by dividing texture space into equally sized cells and measuring distances to “feature points” scattered inside each cell.

Patterns such as the branching and interlocking veins in a leaf can be created by extending the cell-basis technique to use “feature line segments” that connect cells arranged with an implied parent-to-child hierarchy. Assigning an arbitrary point in texture space to its associated cell, a cell-basis function can be created to measure the distance to one or more “feature line segments” that define a complex branching tree structure. The distance to the nearest feature-line segment can create a smoothly varying function that can render a branching texture pattern.

Cell-basis functions identify each cell with a unique point, which is used as a seed in a pseudo-random number generator. The partitioning scheme that defines the cells can be conceptualized as a mapping from points in a region of space to a seed point. During rendering, a sample point is mapped to a cell where the position and orientation of feature geometry contained in the cell can be determined by creating random numbers from the cell’s seed value. For a branching structure to form, the ends of feature-line segments contained in neighboring cells may have to be checked so that a continuous chain of line segments can be formed. Such checking can be done locally using only adjacent cells.

Since the start and end of feature-line segments are known, segments can be given a parameterization that allows features such as tapering towards the top. The described method can also be extended to 3D models where feature geometry contained in cells provides a database of bounded geometry that can be rendered. Level of detail in procedural models can be handled by switching from geometry to surface shaders that use the same cell-basis function to describe the model. Statistical sampling and approximation of the model can be done based on the space partitioning scheme.

By using different partitioning schemes and feature geometry, complex patterns can be assembled without precalculation or simulation.⁴ Shading languages³ with limitations regarding how programs can be constructed can still be used to implement such basis functions as shown in Figure 1, where a zoomed view of branching and interlocking leaf veins is depicted.¹ The shader narrows the size of the cell down to the veins surrounding a sample point. At each step, a different level of branching is created.

References

1. Meinhardt, H. (1998). *The algorithmic beauty of sea shells*. Springer: New York.
2. Peachey, D. (1985). Solid texturing of complex surfaces. In *Computer Graphics, Proc. SIGGRAPH 85*, 279-286.
3. Pixar, *The RenderMan Interface*, Version 3.1, Richmond, CA, September 1989.
4. Prusinkiewicz, P., & Lindenmayer, A. (1996). *The algorithmic beauty of plants*. Springer: New York.
5. Worley, S. (1996). A Cellular Texture Basis Function. In *Computer Graphics, Proc. SIGGRAPH 96*, 291-294.




Figure 1: Branching and interlocking veins as seen in leaves.

Procedural textures are usually computed on the fly as the object is rendered. They avoid the space required by pre-stored images and can define texture more efficiently for solid texturing. However, procedural texturing is an expensive technique, and executing it on a per-pixel basis is costly, preventing it from being used in real-time applications.

The Perlin noise function¹ enhances the realism of procedural textures. A common implementation of this function, called value noise, assigns random numbers to the vertices of a lattice, where all three solid texture coordinates are integers. The noise value for a non-lattice point is obtained by trilinear interpolation of the eight corner points of the lattice cell containing it. While noise is usually the most common function in procedural textures, it is also usually the most expensive. Hence, our desire to perform real-time procedural texturing² leads us to seek fast techniques for computing this function.

To this end, we developed a forward-differencing implementation of the noise function. Forward differencing is familiar to the graphics community, as it is the technique behind the Bresenham line algorithm.³ The value of a pixel is computed from the previous pixel value by adding forward differences. Since trilinear interpolation results in a cubic polynomial, three levels of forward differences are used to calculate the noise value. The forward differencing needs to be set up when the next pixel corresponds to a new cell or the slope of the texture coordinate along the scan line is changed.

noise function can be represented as:

$$\text{noise}(i) = ai^3 + bi^2 + ci + d$$

where i is the distance from the first pixel on the scan line.

The initial value of the three forward differences are:

$$\Delta n = a + b + c, \Delta^2 n = 6a + 2b, \Delta^3 n = 6a$$

At each iteration, noise and forward differences are updated as:

$$\text{noise}_+ = \Delta n, \Delta_+ = \Delta^2 n, \Delta^2_{n+} = \Delta^3 n$$

Rendering time for a marble-textured cube of size 256 by 256 pixels was measured. The algorithm was prototyped in C and tested on a 500 Mhz Pentium III. The average rendering time was 0.20 seconds per frame, and this is currently about 10% faster than the per-pixel computation, though we expect that margin to grow as we optimize our implementation.

A screen shot of animated marble texture is shown in Figure 1.

References

1. Perlin, K. (1985). An image synthesizer. *Computer Graphics (Proc. SIGGRAPH 85)*.
2. Hart, J. C., Carr, N., Kameya, M., Tibbitts, S. A., & Coleman, T. J. (1999, August). Antialiased parameterized solid texturing simplified for consumer-level hardware implementation. *Proc. Eurographics/SIGGRAPH Graphics Hardware Workshop*.
3. Bresenham, J. (1977). A linear algorithm for incremental digital display of circular arcs. *Communications of the ACM*, 20 (2).




Figure 1: Animated marble texture

In “Galaxy Quest,” Dean Parisot’s space adventure, our hero, Tim Allen, is trapped on a hostile planet and being chased by a large and unfriendly creature made entirely out of rocks. Rocks we can do. Creatures we can do. But how do you animate a monster that is supposed to act like a living being and still maintain the characteristics of a rock?

Initial attempts at creating this creature ended up looking like a lumpy Michelin Man and walking like a drunken sailor. Seasoned effects supervisor Bill George and art director Alex Jaeger designed a trimmed down, simpler creature, a little less anthropomorphic and a lot more rock-like. Sunny Wei, our hard surface modeler, took over from there and sculpted rocks and boulders to match the desert environment where the live-action plates were shot. Then lead Rock Monster animator Huck Wirtz began to design an animation style that would give this creature life. One of the things we wanted to avoid was a Rock Monster that looks and moves like a man in a “rock” suit. Huck created an ambling, rolling gait suitable for a twenty-foot-tall menacing creature. Still missing was a naturalistic way for the rocks to interact with each other and with the forces of nature, like gravity and momentum. In order to make the monster look less like an animated armature and more like an angry pile of boulders, a procedure had to be created that would mimic those natural forces. The skills of R&D principal engineer Jim Hourihan, technical animation supervisor James Tooley, and effects supervisor Ed Kramer were brought in to design a technique that would make the rocks act like real moving rocks. Now we had a creature that seemed to have its own intelligence and yet moved with a convincing rock-like movement.

The original plate, a live-action desert environment, was very dusty, with a crumbly pastry-like surface, which meant our Rock Monster needed to be dirtied up to match the look. A method was created to allow the creature to leak debris in the form of dust and small rocks from points of contact on the model. These elements needed to be directly affected by the pressure of the rocks on each other, as well as the momentum with which each rock was moving. Practical dust elements were also added to enhance the effect.

Finally, a creature of this size and weight could not walk through an environment without having some kind of effect on it. Footprints were added, as well the occasional tumbling boulder generated by the creature’s movement.

This may sound like a technique was set up, a button was pressed, and we all sat back and watched it happen. In truth, each scene had to be polished right down to the individual rock in order to get something both aesthetically pleasing and technically correct. To demonstrate all of this, we show a compilation of scenes in which the Rock Monster is chasing Tim Allen, then select one sequence and break it down from start to finish.

Creating a digital human has always been one of the toughest challenges facing computer graphics artists. As people, we subconsciously spend most of our lives paying attention to the nuances of motion and behavior of the people we interact with. This makes us especially attentive to the shortcomings of the digital actors that we see on-screen. On “The Hollow Man” project, this challenge was compounded by the fact that we not only had to create a convincing digital simulation of a specific actor, but this digital human was also expected to be anatomically correct, both inside and out. And while a viewer does not need to be a medical expert to notice when a digital character is anatomically or physiologically inaccurate, extensive research and painstaking attention to detail is indeed required when attempting to recreate such precision. If the modeling, animation, and deformations of a CG character are not anatomically and physiologically precise, the human eye will pick up on this immediately, and the illusion of reality will be broken.

COLLECTING DATA

Before construction of the specific digital actor could begin, it was necessary to gather as much information as possible from the live actor himself. Numerous cyber scans were generated. These included scans of the full body, hand close-ups, and various facial expressions. Because of the potential technical inaccuracies when dealing with cyber data, it was necessary to supplement this information with a large volume of photographs and videos taken from multiple angles and distances. In addition, rulers and calipers were used to take detailed proportion measurements from the live actor. The accuracy of the proportions and posture of the underlying joint structure would be an essential element in successfully replicating the actor’s specific motion characteristics. Therefore, expert medical supervision was always close at hand during the reference-gathering process. Procedures such as X-rays and CAT scans were considered, but ultimately dismissed in order to avoid unnecessary stress and health risks to the actors.

BUILDING THE DIGITAL SKELETON

The technical setup of the joint structure is the foundation for the entire structure of the digital character and can be considered the first step while setting up the animation controls. Accuracy in joint placement is essential for creating body masses that move realistically. The offset of a pivot can cause a limb to appear broken and will often necessitate a reworking of the skeleton at a later date.

PRACTICAL (PRODUCTION) CONSIDERATIONS

It was always important to remember the final goal of creating a useable skeleton. One of the challenges we faced while setting up realistic anatomy is deciding the optimal method for representing the

nuances of a real skeletal structure with the limitations of the current design of digital joints. Joints like the shoulder, wrist, and ankle contain multiple pivot points for each axis. Certain joints shift their rotations between two pivots along the same axis. Traditionally these features are ignored, and the axes are combined onto one joint. For our project, we had a number of internal muscles, which relied on the correct functionality of the internal bones, so it was necessary to include these types of details.

In the wrist, there is a natural distribution of rotation throughout the bone structure. By using these natural bone movements, we were able to thread the muscles through the carpal tunnel. In some cases, like the vertebrae in the spine, we did separate the rotations of each bone between a different joint for x, y, and z. In general, important skeletal characteristics must be left in to maintain believability. It is important to distribute shoulder rotations between the shoulder and the clavicle, and to add details like the motion of the jaw forward while opening. These are the details that maintain realism throughout the character.

CONTROLLING THE CHARACTER

The final step was to build a set of “puppet” controls, which would allow the animators to effectively pose our digital actor. Two very important conditions had to be satisfied in this step:

1. The control set had to be simple enough to allow the animators to work productively
2. The motion of the underlying hierarchy had to maintain all the complexities of realistic human anatomy and motion despite being driven by a relatively simple set of controls.

These two conditions were very often at odds with one another when we created our control system. As sit moves, the human body makes constant natural compensations in order to maintain balance and conserve energy. And most human motion is a great deal more complicated than it looks. Shoulder rotation is virtually impossible without the contribution of the clavicle. Even the slightest pelvis motion involves compensatory rotation of the entire spine. Wrist rotations stem from different pivot points depending on which axis is turning. A complex system of constraints, expressions, and phantom joint hierarchies was necessary to allow for such details to be effectively driven by our reasonably simple set of controls. And in the end, despite our exacting attention to anatomical precision, a small variety of “puppet cheats” were ultimately implemented in order to exaggerate certain motions in the interests of “cinematic license.”

The authors would like to acknowledge Walt Hyneman.

Building “Shrek’s” Natural Environment (Trees, Plants, and Grass)

Scott Peterson
Jonathan Gibbs
PDI/DreamWorks

A challenge in creating the fairy-tale world of “Shrek” was designing large-scale outdoor environments that followed guidelines created by our art directors. Trees and other plants play important visual roles. This sketch discusses some of these issues and some of the technical solutions we are currently using in production.

One goal in “Shrek” was to see objects as simple shapes from a distance, but to still see them as natural and complex when close up. We take two different approaches to this. When designing trees, we start with complex natural branching patterns and then prune them to more simple forms. Other parts of the environment are simple shapes covered with complex arrangements of grass and plants.

A tree system allows animators to grow an entire tree procedurally or grow branches on a hand-modeled trunk. The procedural growth system is an extension of some of the techniques discussed by Weber and Penn.¹ Once a tree is grown, the animator can prune it. This system is written in PDI’s proprietary scripting language, which is also used to procedurally animate trees after they are placed into shots. The trees are saved as a hierarchical collection of splines that represent branches along with several procedural controls for determining leaf placement.

The system for growing grass is the same system used for fur and short hair. We view grass and fur as surfacing operations, so this system is entirely based on image maps and procedural controls accessible in a shader at render time. In addition to color maps for the ground, animators can paint density, length, width, and direction maps to indicate how and where the grass will grow. For other small plants, the modeling department creates models that are then deformed along curves generated in a similar fashion.

Trees, grass, and plants must be able to react to forces such as breezes and actions of the characters. Grass and plants are moved by supplying a force vector to each plant base and tip. Trees are moved by specifying one force vector per branch. For wind, a modified version of PDI’s proprietary fluid simulation system generates wind vectors, and a small simulation models the wind in an environment. The resulting fluid flow moves the plants. A separate system tracks reaction to character motion by tracking character parts, such as feet, and emits a repelling force. Scripts can control force vectors for specialized motion.

A procedural primitive inside PDI’s proprietary renderer generates all geometry. For trees, this system turns the hierarchical splines into a suitable surface and smoothly stitches separate branches where possible. Suitable texture coordinates are generated so trees can be hand-painted. Of course, individually painting 10,000 branches wouldn’t be practical, so branches can be mapped in groups.

All the deformation of grasses and smaller plants, as well as the generation of leaves, is done inside a procedural primitive at render time. This gives us many of the benefits associated with instancing while allowing each object to move and deform independently.

Reference

1. Weber, J. Penn, J. (1995). Creation and rendering of realistic trees. *Proc. SIGGRAPH 95*.

Much of the time required for developing a virtual environment (VE) application is spent modeling objects.¹ Three closely linked attributes of virtual world objects are their appearance, geometry, and behaviour. For an object, the appearance is given by visual properties such as colour or texture, the geometry is the spatial structure of the object, and the behaviour describes how the object evolves or responds to user events. In the design of objects, there are direct trade-offs among these three attributes that affect usability. For example, computational costs to manipulate complex objects (consisting of large numbers of geometric primitives) can be reduced by rendering objects as texture maps and not as individual primitives. Thus, behaviour may require distinct geometric structures. Dually, replacing structures by textures constrains the flexibility for possible behaviour specification. Also, realistic rendering (via texture mapping) may suggest that an object has associated complex behaviour that may be difficult to implement without access to the necessary geometric primitives.

In the development of virtual environments, an object's appearance is commonly defined via a 3D modeler and imported into the virtual world, while behaviour is coded at a low level directly into a prototype application. Unfortunately, finding a suitable level of geometric decomposition for objects in a particular application is a non-trivial problem. For example, if an object's geometric granularity is prematurely completed, then the modeled object may not effectively support the behaviour needed for a particular task. What we seek are methods that will help the designer construct virtual-world objects that will support the user's tasks.

Work to develop guidelines for the specification of virtual world objects for VEs^{3,4} is currently being carried out as part of the INQUISITIVE project.² This is a three-year research project funded by the UK Engineering and Physical Sciences Research Council between groups at The University of York and the CLRC Rutherford Appleton Laboratory (RAL). The project's aim is to develop methods and principles that can be used to improve the design of interfaces for virtual environments.

We propose development of a tool to aid specification of virtual world objects. The method we are developing is based on a standard scene graph, annotated to support storyboarding of interactions in a virtual environment. Our graphs provide a structure to describe the object hierarchy in a virtual environment, and the storyboards allow us to document the behaviour of the environment. We are investigating tool support to document the three components of virtual world objects as part of a structured design process. Specifically, we wish to identify:

- Appearance: Each object is given a rating (background, contextual, or task) based on its role in the user tasks supported by the VE. The rating will be used to decide the level of realism with which the object is rendered.

- Geometry: The required geometric decomposition of objects (in the environment and to object sub-components) to support the behavioural requirements.
- Behaviour: Identification of behaviour in the environment including user/environment interaction, user/object interaction, and object/object interaction.

Figure 1 shows an outline of an annotated scene graph for part of a virtual kitchen interaction. It highlights the identification of object decomposition and requirements for appearance ({} and <>) based on tasks completed by the user.

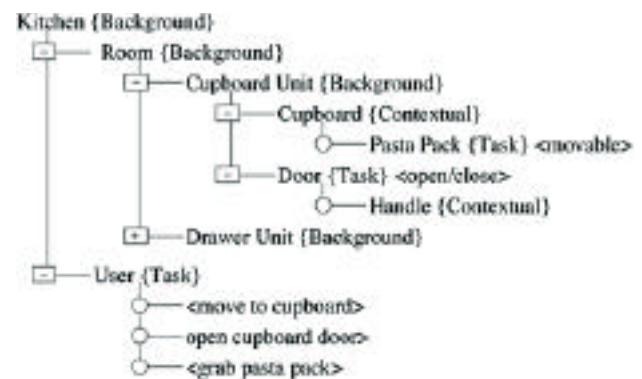


Figure 1: Annotated scene graph example.

At present we are investigating ways of supporting user-defined annotation to capture semantic constraints within the application domain. A tool to support construction of these graphs is currently under design. It will be implemented in Java and made available via the WWW. Our ambition is to integrate this with tool support for prototyping and implementing interaction techniques,⁵ and ultimately with modeling tools. This is work in progress.

References

1. Grinstein, G.G. & Southard, D.A. (1996). Rapid modeling and design in virtual environments. *Presence*, 5 (1), 146-158.
2. INQUISITIVE Project: www.cs.york.ac.uk/hci/inquisitive/
3. Smith, S.P. & Duke, D.J. (2000). Binding virtual environments to toolkit capabilities. *Computer Graphics Forum*, 19 (3).
4. Smith, S.P., Duke, D.J. & Willans, J.S. (2000). Designing world objects for usable virtual environments. Accepted for publication at the workshop on Design, Specification and Verification of Interactive Systems DSV-IS'2000, Limerick, Ireland.
5. Willans, J.S. & Harrison, M.D. (2000). A "plug and play" approach to testing virtual environment interaction techniques. Accepted for publication at the 6th Eurographics Workshop on Virtual Environments (EGVE'2000), Amsterdam, The Netherlands.

Bunshin: A Believable Avatar Surrogate for Both Scripted and On-the-Fly Pen-Based Control in a Presentation Environment

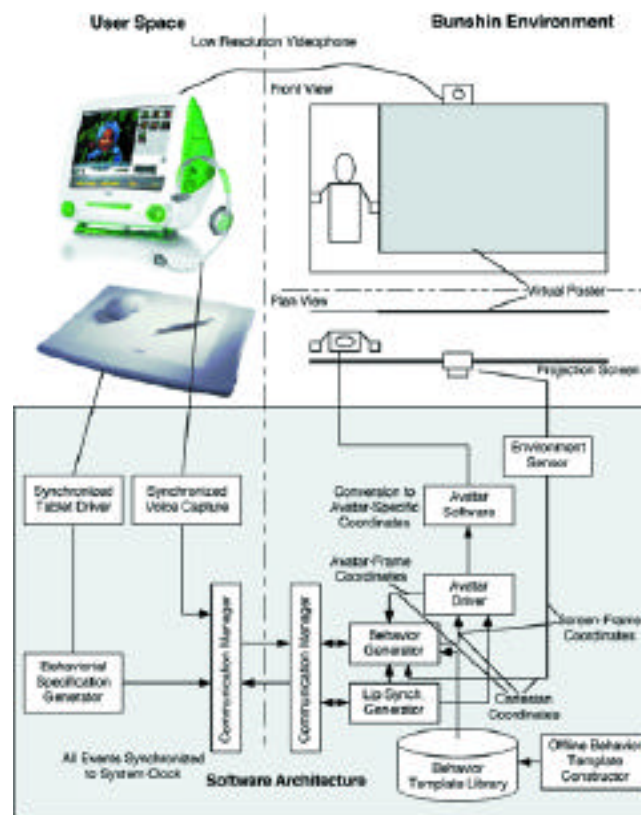
Francis Quek
Wright State University
quek@cs.wright.edu

Richard Yarger
Yonca Haciahmetoglu
Wright State University

Jun Ohya
Kawato Shinjiro
Ryohei Nakatsu
ATR Media Integration &
Research Laboratories


In natural conversation between humans, gesture and speech function together as a “co-expressive” whole, providing one’s interlocutor access to semantic content of the speech act. Psycholinguistic evidence has established that gesture and speech proceed together from the same “idea units,” and at some point bifurcate to the different motor systems that control movement and speech. For this reason, human multi-modal communication coheres topically at a level beyond the local syntax structure. While the visual form, magnitude, and trajectories may change across cultures and individual styles, underlying governing principles exist for the study of gesture and speech in discourse. Chief among these is the temporal coherence between the modalities at the level of communicative intent. An understanding of the constants and principles of such speech-gesture cohesion is essential to its application in HCI involving both modalities. Our remote poster presenter is designed to embody these multimodal components in an avatar controlled in real time through a pen-based interface. Our hypothesis is that gesture-speech synchrony is preserved in pen gestures, and that we can reproduce psycholinguistically correct 3D gesture and gaze behavior from the types and locations of pen strokes.

The environment of the presentation (the Bunshin Environment, BE) comprises a graphics workstation, a rear-projection screen, and a camera that can detect people and events. The software systems in the BE are the Behavior Generator, Avatar Driver, Avatar Software, Communication Manager, Lip Synchronizer, and Environment Sensor. At the user site (the User Space, US), the user interacts with the audience in the BE via a microphone, a low-resolution video-phone, and a pressure-sensitive tablet with graphical overlays. The software components in the US comprise the Pen Tablet Driver and Voice Capture modules, a Behavioral Specification Generator, and a Communication Manager. Depending on the kind and location of the pen gesture, the behavioral specification generator produces a time-tagged “behavior unit” (for example, a concrete deixis at a point in the poster) that is transmitted to the BE via the pair of communication managers. The Behavior Generator receives the behavior units, and information from the Environment Sensor and Avatar Driver (avataric proprioception) to produce appropriate gesture activity. The Lip Synchronizer functions by detecting voicing units in the speech signal from the US. The gestural behavior is described as a set of five 30-sample-per-second vectors describing the position of the hands, the orientation of the head, and the position and orientation of the avatar torso. The avatar driver takes these absolute coordinate vector streams and passes them to the appropriate avatar software API calls.



Anatomical Travelogue and Volume Graphics have created a film of a microscopic alien creature whose mission to earth is to study the human being and the root of intelligence. The creature is totally unfamiliar with the human form but has multi-modal radiological, flexible endoscopic vision. The creature can, on the fly, see the human body multi-dimensionally with MRI, spiral CT, PET, and virtual endo-

scopic powers. This mix of volume data and bio-animated polygons extrapolated from the segmented tissue allows the creature to view the body with no inhibitions. In the end of this baffling journey, during its final fly around and through the body, we begin to realize that the creature is incapable of differentiating human organs, so it attributes intelligence not to the brain but to the colon.



“City of Brass” is a digital film that incorporates both optical and computer environments. Although this process has been used in many recent films, there are a number of novel technical and aesthetic implementations in “City of Brass” that expand the language of cinematic storytelling. This process facilitates the transition from photographic realism to aesthetic impressionism and thus creates a tension between psychological realism and symbolism. The viewer confronts an enigmatic image that draws simultaneously upon the perceptual powers of both the conscious and unconscious mind. The former deciphers the real environment, while the latter is engrossed in the symbolic image. Perception thus encounters a double reality, the familiar and of the symbolic. The ensuing multi-layered, digital image becomes a hieroglyphic sign with a hidden meaning, symbolical and emblematic. The filmic elements are so profusely intermingled with the formal architecture of the digital image that form and meaning become one.

SOCIAL CONTEXT AND FACTORS

Because of recent developments in computer technology and the social impact computers in general have on our culture, “City of Brass” (loosely based on an obscure story of the same title in “The Arabian Nights”) attempts to visualize and reflect upon the social and aesthetic factors computers play in both the East and the West. The film integrates optical and computer landscapes to expand narrative film, but also to meditate upon East/West conflicts. Set in 18th century France, the optical images have horizontal perspectives and vivid colors reflecting the stability and power the West enjoys. On the other hand, computer landscapes mixed with optical images are created in 10th century Arabia with distorted spaces, oblique perspectives, and monotone colors to allegorically reflect the anguish and desperation of the Arab landscape. The fusion of digital and optical landscapes therefore creates a tension of correspondences and dissimilarities. The resulting landscapes dramatize the narrative and create a form of punctuation, which is both structural and thematic. This new form of dramatic mise-en-scene defined by a digital/optical use of space, as opposed to the purely optical, creates detached feelings of sadness, dread, suspense, nostalgia, hopelessness, and lost time.

NARRATIVE CONSTRUCTION

“City of Brass” is the perfect narrative for the digital cinematic image. It has a unique structure in which multi-layered narratives intersect at various points; reflect one another through motifs (color and sound); interconnect the world of kings with that of ordinary people, animals,

and the supernatural to create parables; and above all shift between the first-person narrator, Shaharazad, and the third-person narrators. The formal architecture resulting from juxtaposing optical and computer environments in the film is thus more potent, as the thematic narrative it embodies follows the logic of its form. It is here that form and meaning become undifferentiated. Form and meaning are further intertwined in the presence of a protagonist, Antoine Galland, the famous 18th century French translator of “The Arabian Nights,” who is removed from the actual narrative and acts as a distant observer but who is unable to resist the seduction of the words he is translating. Delirious, jealous, and on the verge of madness, he lives simultaneously in the familiar world of his house and adulteress wife, and in the fantastic world of the story he translates. Antoine’s simultaneous presence in both a realistic world and a symbolic one forces the viewer’s perception to oscillate between two states of mind.

MODELS AND ANIMATION

The animated figures are built after the sculptures of Alberto Giacometti. They are extremely elongated with rough, craggy surfaces. Almost as thin as strings, seemingly weightless and massless, tall, mysterious, and inscrutable, they nevertheless seem to possess and stimulate the surrounding space. They suggest figures long buried underground and now unearthed to stand out in the light. They come into being as if they have suddenly risen out of nothing, as an unalterable whole, and impose on their environment a scared atmosphere.

The figures are intended not merely to describe a real person and be seen, but to draw the viewer’s attention to the act of seeing itself. Their lack of eyes, staccato motion, and unbalanced proportions force viewers to wander over them and around them and reconstruct them again in their own minds. In this way, viewers are directly involved and forced to see them as hieroglyphic signs of a much larger reality, the reality of a mythical content which gives reality itself its vitality.

CONCLUSION

“City of Brass” shows that the interplay of optical and computer environments in digital cinema and their integration into the overall architecture of filmic elements can expand the language of cinema, yielding a cinematic image that carries symbolic and social overtones that require the intervention of the viewer to decipher.

www.evl.uic.edu/arabian-nights/



The evolution of character animation technology has reached the point where additional realism and more richly caricatured behaviors can be amplified by the provision of dynamic clothing on figures. It is now technically practical for clothing to appear on a widespread basis throughout the entire cast of a CG production. Principal characters, crowds, and incidental occurrences of fabric objects in the environment all work together to contribute to a more sophisticated, attractive, and physically familiar environment.

The intention of having a Wardrobe Department (an offshoot of the FX group) for a production like “Shrek” required a significant amount of preparation. It was essential that dynamic clothing processes not become a computational or production bottleneck. Additionally, garments could not inhibit the directors and motion animators from impairing any character behaviors appropriate to the story.

Another requirement of Wardrobe’s clothing system was that it needed to be flexible: It should allow processing of a variety of garments, for a wide range of anatomies, and for arbitrary populations. A character may have one or more dynamic garments (for example, a skirt, a hood, and a cape), each of which might need geometric detail and ornamentation that is character-specific. Every occurrence of simulated fabric might need to exhibit independent physical properties, and these might require individual finessing after the character motion work is completed.

Conversely, one of the challenges facing the Character TD Department was development of a non-dynamic clothing system. The “digital underwear” concept was created to provide a common layer of deformation for characters that were wearing several types of clothes. Procedurally driven cloth elements were able to slide and avoid collisions with skin and other garments underneath, without the need for an expensive simulation. Wrinkles were procedurally created yet still creatively controlled by the animator or technical director.

Prior to production, the design of the dynamic clothing system required collaborative planning with the Modeling and Character TD crews that defined each character’s structure and attributes. The motion animators, who manipulated each character in PDI’s animation environment, were not to feel the impact of the clothing system once production began.

It was necessary to prepare tools that allowed poses and dynamic conditions for garments to be sampled from existing simulations and transferred into new shots for creative or continuity purposes. Also important was the need to ensure that clothing operations utilized the render farm effectively, whether all the garments in a shot were simulating, or only a subset.

Several modules of software were developed that can negotiate all of the stages through which garments must pass, particularly in crowd situations. Ultimately, any complex character-based dynamic processes can be performed on a large scale with timely and predictable results.

In this sketch, we describe a prototype of a novel hybrid 3D object modeling system, NIME - NAIST Immersive Modeling Environment, that inherits the advantages of both traditional 2D GUI-based modeling and 3D immersive modeling environments. Many CAD/CAM systems use 2D GUIs for the modeling interface. There are several systems called “immersive modeler” that enable users to design 3D objects totally in 3D environments.^{1,2} The basic concept of our system is to combine advantages of these 2D and 3D modeling environments in one environment. Forsberg et al. Built a similar system, but it does not allow users to perform modeling in a 3D environment.³

By employing a slanted rear-projection display, NIME integrates 2D and 3D modeling environments into a unified modeling space. By using the surface of the display as a drafting board, NIME presents a 2D GUI modeling interface to the user. NIME also provides a 3D modeling environment by displaying field-sequential stereoscopic images of objects and using a 6-DOF pen-type input device (Figure 2). A user can create 3D models by seamlessly switching between these two modeling environments (Figure 1). The 6-DOF pen-type input device is a combination of a light pen, an inertial sensor, and an ultrasonic sensor. This device is used in both 2D and 3D modeling environments. By calculating the distance between the display surface and the tip of the pen-type input device, the system detects which modeling environment a user intends to use.

A user can perform a number of modeling operations in either 2D or 3D environments, according to the nature of each operation. To create objects, first a user makes a 2D plane shape in the 2D modeling environment. After creating 2D shape, a user can extrude the shape directly into 3D space. We call this way of creating objects “direct extrude.”

After the shape is extruded, the user can edit the intersection of the object. Three additional types of shape creation are prepared: path-extruded, revolving, and pyramidal shaping.

In order to deform an object, NIME provides four types of operations for a user: “intersection extrusion,” “point edit,” “virtual magnet,” and “Boolean operation.” When an object is arranged so that it intersects with the display surface, users can extrude the intersection as they translate the upper half of the object in 3D space, by pushing an “intersection extrusion” button on the screen. At the same time, various 2D operations can be applied to the intersection. In “point edit” mode, a user can move a selected point of an object in both 2D and 3D modeling environments. The “virtual magnet” enables users to deform an object smoothly in the 3D environment. The “Boolean operation” enables users to combine two objects by performing a Boolean operation on solid models of the objects.

By using these modeling operations, users can seamlessly model 3D objects in the NIME system. Figure 3 is an example of modeling an apple by using a 2D drawing, revolving, virtual magnet, and Boolean operation.

References

1. Butterworth, J., Davidson, A., Hensch, S. & Olano, T.M. (1992). 3DM: A three dimensional modeler using a head-mounted display. *Proc. ACM Sympo. on Interactive 3D Graphics*, 135-139.
2. Deering, M.F. (1996). The Holographic VR sketching system. *Communications of the ACM*, 39, (5), 54-61.
3. Forsberg, A.S., LaViola Jr, J.J., & Zeleznik, R.C. (1998). ErgoDesk: A framework for two- and three-dimensional interaction at the ActiveDesk. *Proc. 2nd International Immersive Projection Technology Workshop, IPT'98*.



Figure 1: An overview of the system.



Figure 2: A pen-type input device.




Figure 3: Modeling an apple.



We are concerned with texture mapping recovered geometry with photographs. The geometry is represented as a triangular mesh that may be recovered either from photographs or from laser range images.³ From photographs that have been aligned with the geometry, we wish to manufacture specialized texture maps, images broken up into triangular patches. Since each triangle in a mesh may be covered by multiple photographs, we actually synthesize one triangular texture patch for each triangle to remove the redundancy. This texture patch is an appropriately weighted average of the projected areas of the 3D triangle in all photographs. We use the scheme described by Soucy, Godin, & Rioux¹ to place the synthetic texture patches into texture maps, and therefore obtain texture coordinates.

The size of each texture patch is determined by the maximum number of pixels mapped onto it from any photograph. This approach generates texture patches that are too large. Graphics hardware has a limited amount of texture memory. If we cannot pack all texture information into that amount of space, we need to swap texture memory multiple times for each frame. Swapping texture memory is quite expensive (our SGI O2 needs 0.04 second to load a 1024x512 texture map).

TEXTURE PATCH RESIZING AND CLUSTERING

The amount of memory needed for texture can be reduced by considering the amount of variation in the photographs. We need fewer texture-map pixels (texels) to represent smooth areas than to encode highly varying regions. We use the response of an edge-detection operator (the derivative of the Gaussian) as our information measure and apply it to all original photographs. For each texture patch, we use the maximum response at its corresponding pixels in the photographs to determine the number of texels it actually needs to keep the original color variations on the triangle. Thus smooth regions where the operator does not respond are allocated few texels.

Additional savings can be achieved by reusing texture patches for multiple 3D triangles, when the texture over these triangles looks similar. This requires that we cluster the texture patches and set the same texture coordinates to all triangles in the same cluster. We force the number of texels along different edges of a texture patch to be the same and quantize this number to be a power of 2. Then we use the K-mean algorithm (Lloyd algorithm) in vector quantization² to cluster all the texture patches with the same size. Because of Mach Band effect, slight color differences along the edge shared by two 3D triangles may be rather obvious. We use a larger penalty for difference on edge texels to alleviate this effect during K-mean clustering. Given an error tolerance, we need to run a binary search to find the minimum number of clusters that can achieve that error. This process is quite time-consuming since each step of the binary search needs to run the K-mean algorithm whose complexity is $O(nmd)$ where n is the number of initial vectors, m is the number of clusters, and d is the dimensionality of each vector. We found out that a two-level scheme can optimize the performance by first grouping the n vectors into square-root(n) clusters and then running the binary search on the vectors belonging to each cluster, which in turn splits into multiple clusters.

1. Soucy, M., Godin, G., & Rioux, M. (1997). A texture-mapping approach for the compression of colored 3D triangulation. *Visual Computer*, 12, 503-514.
2. Gersho, A. & Gray, R.M. (1992). *Vector quantization and signal compression*. Boston: Kluwer Academic Publishers.
3. Yu, Y. (2000). *Modeling and editing real scenes with image-based techniques*. Doctoral thesis, Computer Science Division, UC Berkeley, 2000 (www.cs.berkeley.edu/~yyz).





This image shows a real room rendered with 20 texture maps.



This second image shows the same room rendered with only five texture maps produced from clustering.

The Anatomical Travelogue, Volume Graphics, the University of Bologna, and the Armed Forces Institute of Pathology are working collaboratively to bring a four-hour television series to National Geographic on the development of the fetus from “Conception to Birth.” Combining their collective talents and technologies, these groups are, for the first time, scanning high-resolution datasets of the male and female reproductive systems down to a single sperm and ovum.

Using 3D visualization techniques and computer animation, we peer inside a human egg, a luminous orb packed with its genetic load of 23 chromosomes encased in a protective membrane, enabling us to examine real human data in ways not possible with any other technology. Using these same technologies, we join an invading army of sperm as they swim upstream to start the process of life, witnessing the extraordinary encounter of sperm and egg, capturing the precise moment of conception. We witness the transformation of a single fertilized cell to differentiating series of cell divisions, isolating every wondrous step in this extraordinary evolutionary process, the most important developmental stage of human life. Using real medical data to visualize the once invisible, untold biological mysteries are unraveled.



Since Newton and Louis-Bertrand Castel in the 17th century, scientists and artists have been interested in the relationship of light and sound. And from Scriabin, who wrote a light score for Prometheus, through David Bowie, who observed that “the eye is hungrier than the ear,” composers have recognized the potential of the visual to enhance our enjoyment of music. Over the last century, painters and sculptors who aspired to the dynamism and immediacy of music invented light organs, created kinetic sculptures, and produced animated films.

Computer graphics and MIDI have made it possible to integrate music and images as never before. Imager, a computer-based instrument, is designed to let artists play graphics in the way that musicians play with sound. Its design relies on developments in modern art to provide control of color, form, and movement for use in both composition and improvisation.



Imager’s color model supports, and even encourages, application of hundreds of rules of thumb and principles that have been developed by painters and color theorists. These rules often rely on concepts such as hue, saturation, value, tint, and shade. The hsv (hue, saturation, value) color model makes representing such concepts straightforward. Further, by providing for conceptual relationships among colors as complementary or adjacent, warm or cool, dark or light, pale or intense, it facilitates changing the colors of large numbers of differently colored objects improvisationally.

Form is composed from simple elements: lines, arcs, polygons, conic sections, and spirals. Because these elements are simple, many can be rendered and updated simultaneously. Together, they form a vocabulary that is familiar to artists trained in the past century or so. And out of them it is possible to devise a diversity of images and image styles that are not easily characterized as coming from a single aesthetic or “instrument.”

While rhythms can result from changing colors and forms, the most important source of visual rhythm is from the movement of the forms. Forms can grow and shrink, rotate around axes, and move along designated paths. The parameters can be altered while listening to musicians play. Max provides a convenient environment in which to specify how various MIDI controllers are used to trigger graphic events and transformations. The amount of learning required to play the instrument seems of the same order as for many musical instruments. It is possible to establish basic rhythms fairly quickly, but more subtle control results from practice and from refining mappings for particular music.

The 20th century opened with explosive scientific developments in our understanding of light and color. Albert Michelson, one of the preeminent researchers in the field, was excited by what he saw: “Indeed, so strongly do these color phenomena appeal to me that I venture to predict that in the not very distant future there may be a color art analogous to the art of sound — a color music.” Perhaps the “not very distant future” is upon us.

imagers.cwru.edu



Realism in computer graphics can be obtained in two ways. In photorealistic rendering, synthetic images are produced that look “real.” Computation of such images usually requires minutes or hours. But realism can also be created through interactivity, even if the world the user explores looks merely synthetic.

The gap between photorealistic and interactive rendering is still large. The rapid development of graphics hardware makes interactivity possible for complex scenes, even on cheap PC graphics boards. Nevertheless, just the inclusion of shadows as the most basic global illumination effect usually destroys interactivity. On the other hand, even high-end multi-processor machines are not fast enough to achieve interactivity using ray tracing for more than trivial scenes. The Holodeck¹ and render cache² systems were big steps for making ray tracing interactive, but they still require enormous computation power and leave the valuable graphics hardware nearly unused.


We try to bring together interactivity and ray tracing by a hybrid approach. In our system, a user moves through a world interactively. Views are generated using graphics hardware, with all known limitations in realism. Whenever computation power is available on local or remote processors, we acquire ray tracing samples in parallel to the rendering of the current view. The resulting samples are then used to generate “corrective textures” that represent the error between the hardware-generated images and the ray-tracing results. These corrective textures are mapped onto the objects during hardware rendering, correcting the hardware-generated image towards the ray-tracing solution.

Several issues have to be addressed by this approach:

- **Assignment of textures:** We do not assign a texture to each scene triangle, but to entire groups that form an object (for example, a cup, a chair, or a book). For this purpose, a scene hierarchy is needed. The hierarchy level to which the textures are assigned depends on the screen size of the object and on its depth range within the current view.
- **Texture mapping:** Corrective textures are projected onto objects by point projection from the current viewpoint. In this way, ambiguities between ray samples and texture pixels are avoided.
- **Insertion of samples:** Samples are splatted into the textures using a Voronoi diagram creation method. Blending is used to blur the Voronoi cell boundaries.
- **Sample generation:** Ray samples are concentrated in regions that exhibit the largest error. The samples can be generated on parallel processors or remote machines.
- **Aging of textures:** Textures are reused from frame to frame, but the splatted samples age, according to the camera movement. Old and thus uncertain samples are likely to be overwritten by new ones.

In our implementation, we render a radiosity solution with graphics hardware and use a standard distribution ray tracer to compute the corrective textures. Even for rather complex test scenes, we achieve display rates of several frames per second.

1. Larson, G.W. (1998). The Holodeck: A parallel ray-caching rendering system. In: *Proc. Second Eurographics Workshop on Parallel Graphics and Visualization*.
2. Walter, B., Drettakis, G., & Parker, S. (1999). Interactive rendering using the render cache. *Rendering Techniques '99 (Proc. 10th Eurographics Rendering Workshop)*. Springer, 27-38.



Snapshots from an interactive session. Upper left: almost converged rendering of a scene after a few seconds. Upper right: new viewpoint, with the same corrective textures. Lower left: after 10.000 new samples have been shot. Lower right: positions of the new samples.

In creating a synthetic Dr. Sebastian Caine for “The Hollow Man,” we were forced to expand creative and technical realms. Choices affected our digital processes and the real-world processes of photography, interaction, and performance. Using our film as a case study, we discuss performance, directing, and shooting considerations as well as the creative and technical challenges of post-production.

Computer animation and imagery have often been seen as a medium that provides individual creative power but is, in fact, most powerful as a collaborative and structured creative process. How do we reconcile this with the inherently personal performance provided by an actor working in conjunction with other actors and a director?

This sketch provides an overview of creative and technical challenges. Additional sessions go into the further detail on the actual applications of modeling, puppet control, performance, and imagery.

CREATIVE OVERVIEW

“The Hollow Man” was a project designed by the writer, the director, and the effects facility to be a state-of-the-art exercise in visual effects. The effects were designed to compliment the story, which was a psychological horror-action-thriller based on an invisible-man theme. The endeavor would take as many as 150 visual-effects artists working over a period of two years to deliver over 400 shots. The work would include many effects that would be unseen by the viewer but were necessary to complete the film’s illusion, as well as the starring role given to the transformation and the climactic scene’s character.

Visual effects are as much visual magic as slavish recreation of reality. Combining technologies is often both the most cost-effective and the most visually satisfying method of solving a problem. Using any tool only when necessary and constantly changing the trick makes it harder to pick up on the sleight-of-hand. Our choices were based on the central tenet of working with as much of the principal photography as possible, thus providing the human cast something to perform to. The technical approach was multidisciplinary, involving cutting-edge research into digital, makeup, and shooting technologies.

We surveyed and analyzed any and all applicable technologies, from the traditional to the digital. While our primary body of work involved creating the digital human, the supporting technologies involved motion control, “rough clean” technology, matte-able clothing, digital cloth, matte-able makeup, and thermal imaging.

In recreating reality, or in creating a performance, one must first study the goal. We drew our studies not only from life sciences, but also from such esoteric sources as the 18th century wax sculptures of “La Specola” in the Natural History Museum of the University of Florence. These waxes provided an important touchstone, showing balance between medical accuracy and creative expression. In addition, we utilized the skills of both artistic and medical experts in physiology. Our research was supported by numerous videos and books on the subjects of physiology and anatomy, and completed by the crew’s involvement in a human dissection.

On one level, our goal was simple: recreate the performance of our lead actor, Kevin Bacon. We had to be very literal and true to Kevin’s performance as a person and an actor. Further, our character would have to interact with the other actors. For this reason, we chose to work with key-frame animation and matching to real or reference performances provided by Kevin himself. Motion capture was deemed too invasive and unreliable in an “on-set” environment, and off-line or studio-based systems eliminated the interaction and on-set performance. Finally, any gains provided by motion capture were soon outweighed by the advantage of having Kevin provide his own “secondary effects.” Splashes, reactions, objects moving, etc. would all have to be recreated with physical or digital effects, had we not used Kevin on set.

Since Sebastian was “invisible,” or at least not completely solid, we had to tackle replacing the background where he previously existed but would no longer be visible. Three technologies were used to provide clean backgrounds: motion control, a technique involving “rough clean” passes, and synthetic recreation (in limited cases).

Putting it back together involved compositing combined with match moving “wild plates” and/or bringing the motion control information into the computer. The match moves were used in the creation of the “missing pieces,” involving the synthetic Sebastian or his world. The mask, sheets, clothing, and other objects were all painstakingly recreated and blended with the photographic elements and Sebastian to create the final illusion.

CONCLUSION

Our work would have been impossible without the assistance, patience, and skill of Kevin Bacon. He provided us with not only his body but, more importantly, his performance. This was crucial to our success.

While many of the technical pitfalls of human animation can and have been solved, the larger issue of performance and direction has not been answered. An actor’s performance, while structured by a director and influenced by his fellow cast members, remains an intensely personal and private act performed either in public or in front of a camera. Our experience of blending the live action and synthetic worlds has shown the power and limitations of both.



Displacement mapping is a common feature found in commercial renderers that has not yet been implemented using hardware in modern 3D graphics rasterization pipelines. Displacement Mapping¹ perturbs a Point P on a surface to a new point P' on the displaced surface using the formula :

$$P'(u, v) = P(u, v) + D(u, v)\hat{N}(u, v)$$

where D is the displacement scalar and \hat{N} is the unit surface normal at P .

Applying a displacement map to a triangle mesh involves retriangulating the original mesh and displacing the vertices accordingly. If the base triangle mesh has a coarser resolution than the triangle mesh, we need to re-tessellate the mesh according to the surface defined by the displacement map. This re-tessellation can be performed while rasterizing the triangle as presented in Doggett & Kugler³ and Gumhold & Hüttner.⁴ The problem with these approaches is that a large number of triangles is generated. To solve these problems, we propose to tessellate the individual triangles sequentially by recursively adding vertices along edges and applying three termination conditions.

EDGE TESTS FOR VERTEX INSERTION

Displacement mapping requires a coarse triangle mesh that approximates the surface to be modeled with a displacement map containing the finer geometric detail. To avoid generating cracks (produced by t-vertices) in the displaced surface and allow processing to be performed one triangle at a time, we limit our vertex insertion decision to individual edges.

Re-meshing is performed by inserting vertices at the midpoint of edges if two conditions dealing with surface normals and displacements are lower than given thresholds. First, the midpoint M between the two vertices V_1, V_2 of an edge is calculated. By bump mapping the original mesh normal using the displacement map normal we can calculate the new normal at M . Then we compare each component of the new normal with the normal components of V_1 and V_2 . If the difference between the components is greater than a set threshold, then a vertex is added at M (**Normal Test**). The Normal Test is subject to aliasing because it uses point sampling and can miss changes in height. A simple example of this is shown in Figure 1(a). To detect average changes in the height of the displacement map, a second test is performed that compares the displacement over an area using Summed-Area Tables.² A summed-height table using the displacement map can be precomputed and the summed height at vertices V_1, V_2 , and M can be calculated from the table by using the point position plus a size determined by the length of the edge. If the difference between the summed height at M and the average of the summed height at V_1 and V_2 is above a set threshold, then a vertex is added at M (**Summed Height Test**). This test misses some cases that the Normal Test detects, as shown in Figure 1(b).

HARDWARE ARCHITECTURE

To achieve hardware rendering of displacement maps, we propose to introduce a Re-meshing Unit into the rendering pipeline, as shown in Figure 2. This Re-meshing Unit generates new triangles from each triangle sent from the CPU before it enters the rasterization pipeline. To calculate the new normals on the displaced mesh, the original mesh




Figure 1: The dashed lines indicate the area over which the height is averaged to calculate the Summed Height value. In (a) the Normal Test fails, but the Summed Height Test succeeds. In (b) the Normal Test succeeds, but the Summed Height Test fails.

normal can be bump mapped using the pre-calculated normal map in hardware using techniques such as those suggested by Peercy, Airey, & Cabral.⁵

RESULTS AND FUTURE WORK

We have implemented this technique in software, and results for the Normal and Summed-Height Tests are shown in Figure 3. The combination of both tests can actually produce a reduced number of triangles due to an improved adaptation to the surface represented by the displacement map. We are investigating the possibilities of saving the mesh state from one displacement map and adding another displacement to it, as shown in Figure 3(d). As displacement maps are increasingly used to model the geometric detail on surfaces, hardware support for rendering these displacement maps will become a sought-after feature in graphics accelerators.

A complex mesh generated from a cylinder is shown in Figure 3. As displacement maps are used increasingly to model the geometric detail on surfaces, hardware support for rendering these displacement maps will become a sought-after feature in graphics accelerators.




Figure 2.



Figure 3.

References

1. Cook, R.L. (1984). Shade trees. *Computer Graphics, Proc. of SIGGRAPH 84*, 18 (3), 223-231.
2. Crow, F.C. (1984). Summed-area tables for texture mapping. *Computer Graphics, Proc. of SIGGRAPH 84*, 18 (3), 207-212.
3. Doggett, M. & Kugler. (1999). *A hardware architecture for displacement mapping using scan conversion*. Technical Report WSI-99-12, Wilhelm-Schickard-Institut für Informatik, Universität Tübingen.
4. Gumhold, S. & Hüttner, T. (1999, August). Multiresolution rendering with displacement mapping. In *Eurographics/SIGGRAPH Workshop on Graphics Hardware*, 55-66.
5. Peercy, M., Airey, J. & Cabral, B. (1997). Efficient bump mapping hardware. In *Computer Graphics, Proc. of SIGGRAPH 97*, 303-306.


Physics-based animation is increasingly being used to create realistic, autonomous animations. Although commercial systems have been developed that produce passive dynamics, lifelike creatures must have their internal forces and torques for locomotion actively controlled. Other simulation systems are applicable to only one type of object representation like articulated figures or deformable cloth.¹

Control schemes have been developed in different customized simulation environments.² To remedy the current fragmentation of simulation systems and control schemes, we have developed a software framework called DANCE (Dynamic Animation and Control Environment). DANCE provides a mechanism to easily integrate different controller schemes and simulate diverse object representations together. With this environment, new practitioners can quickly focus on new techniques without investing large amounts of time to build support infrastructure.


BASE CLASSES

We identify a minimal set of object-oriented abstract base classes that capture many features common in existing simulation systems. The small number of different classes offers a clear perspective of the components of a physical system. Specialized features are added by creating subclasses of each base class through dynamically loaded plug-ins. A central driver orchestrates the communication and management of the plug-ins.

A system describes an entity that can undergo physical change of state. In general, the system class can represent any physical system that can be described with a set of generalized coordinates. We have implemented articulated figures, deformable parametric surfaces, and particle systems. The geometry class provides graphical representations for systems.



A physics-based human responding to pushes and performing a stunt.



An anatomically based modeler that allows muscle, ligament, and skin simulations.

Actuators are entities that exert forces or torques on systems. Gravitational fields, rigid body collision, and biomechanical muscles can be modeled as actuators. Actuators can manage other actuators to form complex hierarchical controllers.

Simulators update the state of systems over time, using either physics-based or kinematic processes. In DANCE, the degrees of freedom of all systems can be handled by different simulators. Conversely, simulators can be assigned to individual systems or groups of systems.

DANCE employs application programming interfaces (APIs) that are available on all major systems: OpenGL for 3D graphics, Tcl/Tk for the scripting and graphical user interface, and GLUT for window management and input event handling. It has been ported to Windows NT, Linux, and Irix. The GUI is completely external to the DANCE system. Therefore, different applications can have customized interfaces.

APPLICATIONS

We have used DANCE to implement two significantly different research projects: a biomechanical, musculoskeletal simulation system featuring deformable musculotendons and ligaments as actuators and a software control system for composing a diverse array of different motion controllers for autonomous and versatile virtual stuntmen. We believe that DANCE provides a useful environment for prototyping physical systems and controllers, and sharing results.

DANCE is being actively enhanced and is available at:
www.dgp.toronto.edu/DGP/software/dance/dance.html

References

1. Baraff, D. & Witkin. (1998). A. Large steps in cloth simulation. In *Computer Graphics (Proc. SIGGRAPH 98)*, 43-54.
2. Hodgins, J. K., Wooten, W. L., Brogan, D. C., & O'Brien, J. F. (1995). Animating human athletics. In *Computer Graphics (Proc. SIGGRAPH 95)*, 71-78.

Many solid modeling applications require information not only about the geometry of an object but also about its “topology,” the connectivity of its faces, edges, and vertices. Many interchange formats do not provide this information; therefore, the application must derive it as it builds its own topological data structure from an unorganized list of triangles. For very large datasets, the data structure itself can be bigger than core memory, and a naive algorithm for building it can become prohibitively slow due to memory thrashing. In this sketch, we describe a new out-of-core algorithm that can build our topological data structure, LEDS (the Loop Edge-Use Data Structure), which is closely related to Weiler’s radial edge data structure,¹ from very large datasets.

To determine connectivity relationships, we use hash tables to match up all coincident vertex coordinates and edge-uses. One possible approach is to build these hash tables and construct and update the LEDS as the input is read. Unfortunately, for a large file, each new element read could be connected to elements that have already been written out to disk, and those elements will now need to be paged back in to be updated. Even if the input is extremely coherent, so that updates to the same element are closely spaced in time, we will still see thrashing when the hash tables become too large to co-exist in memory. The problem of random update accesses is also encountered in bulk loading object-oriented databases; Wiener’s algorithm in this domain² inspired our approach.

Our algorithm for non-memory-resident data avoids thrashing in two ways: by re-ordering and grouping random hash table accesses, so that we need to build and access only one memory-sized partition of a single larger hash table at a time, and by using external merge-sorts

to re-order the intermediate data generated using the hash tables. This allows us to write all of the information that needs to be recorded in each entity at creation time, and we never need to go back and modify entities that have already been written out to disk. Our only out-of-order accesses are within the in-memory hash table partitions and during the sorting stage.

For testing, we took a simple curved shape and varied the fineness of the triangulation to produce files of different sizes (Figure 1). To extract the effects of input coherence, we made two versions of each file: one with the triangles organized in consecutive triangle strips, and one with the same triangles in random order. Running on an SGI Indy with 32M of RAM, the naive approach is efficient for building a small memory-resident LEDS, with identical performance on random and coherent input. For larger files, however, performance degrades rapidly due to thrashing, particularly for the random files. For small files that fit in memory, the out-of-core algorithm takes approximately 20% longer to build the LEDS than the naive algorithm, but the intelligent use of virtual memory more than makes up for this overhead on larger files (Figure 2). The out-of-core algorithm we present makes building a very large topological data structure feasible, regardless of the coherence of the input.

References

1. Weiler, K. (1988). The radial edge structure: A topological representation for non-manifold geometric boundary modeling. *Geometric Modeling for CAD Applications*, 3-36.
2. Wiener, J.L. & Naughton, J.F. (1995). OODB bulk loading revisited: The partitioned-list approach. *Proc. of the 21st International Conference on Very Large Data Bases*, 30 (4).



Figure 1: Our test part, coarsely tessellated to show triangle-strip organization.




Figure 2: Naive vs. out-of-core build times for the test part triangulated at various resolutions. Note that the results for the out-of-core algorithm on the coherent and randomly ordered versions of the input are coincident.

FORCE FEEDBACK AND ITS CURRENT LIMITATION

Over the past decade, various types of special feedback devices have been regarded as a primary target of R&D in the field of virtual reality. Effective interactions are realized by stimulating human sensory organs in various ways, leading to auditory sensation, force sensation, and visual sensation. However, high cost and lack of standards still make it difficult to incorporate those special interaction devices in regular computer systems. This motivated the authors to achieve a pseudo-haptic representation using commonly used interaction devices, such as a normal 2D display and a mouse. The present work can be viewed as a psychological approach, in that we take advantage of a synesthesia to achieve the pseudo-haptic sensation. This sketch exploits a new style of user interface based on the pseudo force sensation by developing a software library for pseudo-haptically enhanced virtual interaction devices.

INTERACTION WITH PSEUDO FORCE FEEDBACK

The authors devote special attention to synesthesia, a psychological phenomenon in which one type of stimulation evokes the sensation of another. We frequently encounter situations in our computing lives where the motion of an object icon, which is hard to move, makes us feel that the object is heavy (vision + motion > force). This manifests the potential for pseudo force feedback using common I/O devices.

By taking advantage of the above-mentioned synesthesia, the authors have developed a software library for pseudo-haptically enhanced virtual interaction devices. Response latency of virtual interaction devices generates a pseudo feeling of heaviness. Table 1 shows our

library for pseudo-haptically enhanced virtual interaction devices, which can be used for indicating to the user the remarkability and importance of corresponding operations. Our work differs from prior work with hardware invention² only in the need for a normal mouse as an input device, and it differs from software development¹ in that it applies the same principle of pseudo force sensation to the entire list of virtual interaction devices.

CASE STUDY: MAP NAVIGATION

The effectiveness of the present library is illustrated with an application to map navigation. In the map navigation system, a map is panned with scrollbar or mouse dragging. In general, panning is likely to make the user lose track of a landmark. The present library ameliorates the problem. The pseudo-haptic scrollbar becomes heavy around the landmark, or the pseudo-haptic mouse pointer becomes lighter as the mouse approaches the landmark. This enables the user to roughly locate the landmark while panning the map. Future application of the library to various general input devices, such as touchpads or accu-points, is one of the most attractive directions for future research.

References

- Ahlberg, C. & Shneiderman, B. (1994). The Alphalider: A compact and rapid selector. In *Conference Proceedings on Human Factors in Computing Systems (ACM CHI 94)*, 365-371.
- Kelley, A.J., Higuchi, T., & Salcudean, S.E. (1995). An integrated haptographical user interface using a force feedback mouse. In *Proc. ATIP95.65 Human Computer Interface International*, 713-718.



Figure 1: Map navigation system using pseudo-haptic scrollbar and mouse pointer.

Virtual interaction device	The effectiveness of pseudo force sensation
Scrollbar	It is possible to roughly locate the landmark while panning the map.
Mouse pointer	It is possible to roughly locate the landmark while panning the map.
Mouse	It is possible to roughly locate the landmark while panning the map.
Touchpad	It is possible to roughly locate the landmark while panning the map.
Accu-point	It is possible to roughly locate the landmark while panning the map.

Table 1: Software library for pseudo-haptically enhanced virtual interaction devices and its effectiveness.

Current rendering systems are now capable of accurately simulating the physical distribution of light in an environment. However, this physical accuracy does not ensure that the displayed images faithfully represent reality. Can we measure how close to reality these images are perceived to be? Existing measures of image fidelity take account of certain characteristics such as the Fourier content of an image, but this may not reveal specific errors in the portrayal of lighting and shadows; furthermore, it could be argued that one needs a measure of fidelity that involves assessment of performance on a specific visual task, so that performance on that task in a real scene can be compared with performance on that task in a graphical scene to give a measure of fidelity. By conducting such a set of experiments, based on the psychophysics of lightness estimation, we can estimate how close a rendered image is to its original counterpart.³ Does this experimental framework using human observers offer advantages over existing

image-comparison methods? In order to demonstrate the differences between using algorithmic metrics and human evaluation we generated a series of images using different lighting simulation programs, Radiance² and RenderPark.¹ Resulting images were then compared to the real scene they represent, first using metrics (comparing to a captured image of the real scene), then using a psychophysical lightness estimation task in the rendered and real scenes. In this sketch, we present the results of these comparisons.

References

1. www.cs.kuleuven.ac.be/cwis/research/graphics/RENDERPARK/DOC/
2. Larson Ward, G. (1998). *Rendering with RADIANCE: The art and science of lighting simulation*. San Francisco: Morgan Kaufmann.
3. McNamara, A., Chalmers, A., Troscianko, T., & Reinhard, E. (1998). Fidelity of graphics reconstructions: A psychophysical investigation. In *Proc. of the 9th Eurographics Rendering Workshop*, 237-246.

A wide class of operations on images can be performed directly in the wavelet domain. The basic idea is to express the operation using linear combinations of matrices, some of which are related to the image being operated upon, while others represent the operation itself. Given the linear combination of matrices, we compute the wavelet transform for each matrix and perform the combination on the wavelet coefficients. Applying the inverse wavelet transform to the resulting coefficients yields the desired result. This sketch describes applications of this approach to 3D image warping and image convolution.


Operating in the wavelet domain has several advantages. Wavelet decomposition of an image or a large matrix is typically much sparser than the original direct representation, resulting in faster computation of linear combinations of matrices. The wavelet decomposition represents each matrix at multiple scales, which allows operations to be performed at multiple resolutions, as well as in a progressive, coarse-to-fine fashion. The wavelet transform has nice localization properties; therefore, by selectively operating on subsets of coefficients, it is possible to restrict the effect of the operation both spatially and in frequency.

The central computational component of many image-based rendering algorithms is 3D image warping, which maps reference image pixels to their coordinates in a target image. Most of this operation can be expressed as a linear combination of matrices, and carried out in the wavelet domain. The resulting algorithm is as fast as the most efficient warping algorithm known to date for the case of planar and cylindrical reference images, and is faster for spherical reference images. Another major advantage of warping in the wavelet domain is that warping can be performed not only at the resolution of the original reference image,

but also at coarser resolutions. Suppose we have an object represented by several high-resolution reference images. The high resolution may be necessary for a close-up view, but for most views of a 3D scene containing the object, a lower resolution suffices. Our approach makes it possible to perform the warp at the appropriate coarser resolution without unnecessarily warping each and every pixel in the reference images. The advantage over pyramid-based representations is that a low-resolution result can be progressively refined without redundant computations.

Convolution is probably the most widely used operation in image processing. By rewriting the operation as a weighted sum of shifted and transformed images, we obtain wavelet convolution, which is more efficient than standard convolution, even for small images and small kernels. By performing lossy wavelet compression, accuracy can be gracefully traded off for speed, providing a very rapid way to obtain a visually accurate approximation to the result of the convolution. Both exact and approximate wavelet convolution should prove very useful in image editing applications, since they result in faster response time, thus allowing users to preview the results of their operations interactively and at the full resolution of the image. Wavelet convolution can also be performed at different resolutions. Furthermore, by performing the operation on subsets of the wavelet coefficients, convolution can be restricted to operate only on a certain region in the image, or only on features at a certain scale, or both.

Additional information can be found on the SIGGRAPH 2000 Conference Abstracts and Applications CD-ROM and at: www.cs.huji.ac.il/~danix/research/wavops/index.html



The top left image demonstrates localized wavelet convolution restricted to a spatial window and to features at a single scale. The other three images show multi-resolution wavelet convolution applied to multiple scales in a progressive, coarse-to-fine, fashion.

We present a method for computation of Catmull-Clark (CC) subdivision surfaces that fit a cloud of vertices. This is a common problem when scanning a real model with a laser and trying to represent the scanned data set as a simple polygonal model. The reason for using CC surfaces is explained in detail.

ABSTRACT

EMP (Experience Music Project) is a new and exciting Digital Domain project whose goal is to describe a history of music as visual experience. The CG supervisors for this project presented us with a very interesting challenge: given a reduced set of laser-scanned points of a James Brown clay model, compute a subdivision surface that would fit the data set. We were given one week to complete a working prototype and another week for a final production tool. Another limitation was that the interpolating subdivision surface^{1,2} had to be a Catmull-Clark surface, since this was the only one supported by the RenderMan renderer used by EMP. Under no circumstance were we allowed to produce a high-resolution tessellated model of the James Brown model. The whole point of using a small set of data (surface subdivision cage) was to reduce the I/O operations during rendering and animation.

One week to complete a working prototype was one of the most important constraints. Equally important was finding a numerical algorithm simple and robust enough to keep development time down and ensure a successful result. At the same time that the software group started developing a technical solution for this problem, the CG supervisors at EMP maintained a second track just in case we failed to meet our deadline. Digital artists at the EMP group adjusted the polygonal mesh by hand for the reduced scanned model and adjusted the mesh by trial and error until the CC subdivision surfaces looked close to the original clay model.

Our first proposed answer to the problem was to use the Jos Stam³ parameterization of Catmull-Clark surfaces to come up with minimization formulas to fit the CC surface locally to the cage data set. However, we opted not to follow this lead in our first implementation because it would have required major changes in our surface subdivision engine, and private conversations with other researchers led us to think that the convergence of such minimization through the Stam parametric form may suffer from unexpected divergences. We spent a couple of days trying to determine the convergence properties and a simple closed form for the fitting formulas, but we saw no clear answer to both questions. So we put this approach on hold for the moment, even though we felt that in the long run it is not only mathematically more satisfactory, but it would also provide the best performance.

With five days left to complete the first prototype, we decided to approach the problem with a very stable, though numerically less efficient, method. Given an initial mesh, we used nonlinear minimization to find the set of local modifications to the vertices that would yield a CC surface whose points lock onto the initial scanned data within an epsilon (10^{-6} in most tests) error. This is a classic surface fitting problem and required only minor modifications to our existing subdivision libraries. The surfaces obtained through this method were highly accurate, and although some of the calculations for the optimized models took a very long time (one hour for the young James Brown model consisting of several hundred vertices), it was still worth

it. Optimization is a step required only once per model, so one hour of CPU processing time is still nothing compared to the time and trouble of artists fixing the model by hand.

At the two week deadline, we compared our results with the model fixed by the artist. The good news is that the numerical solution gave a highly accurate solution and revealed a number of details in the original clay model that were not as visible in the hand-made model. However, there was a major shortcoming in our numerical solution. During the requirements stage of the project, we neglected to consider other issues besides fitting the best surface to the scanned data, mainly the animation setup requirements for such a model.

When we compared the numerical solution with the artist solution, it became clear that the user was able to adjust the location of the edges in the model along lines that made it easy for an animator to create deformations in the face. Keeping the proper alignment of edges along the lips, cheeks, and other facial expressions was just as important as generating a highly accurate Catmull-Clark representation of the scanned data. Our numerical tool was oblivious to this issue and generated arbitrarily aligned cage edges. For this reason, the final production for the EMP shot used the artist's model instead of our tool.


In conclusion, we have gone through a production cycle that taught us a great deal about the desired requirements for a Catmull-Clark fitting tool. Our experience tells us that the obvious requirements of precision in the fitting of the surface, though important, are by far not sufficient. Equally important, from the user's point of view, is to allow the artist the possibility of outlining lines of constraint for the minimization (or surface-fitting) problem. Our sketch reviews on-going work in a second generation of the surface-fitting tool that includes the above enhancements.

Beau Cameron
Young James Brown model

Giancarlo Lari
Bat model, Rhino model

References

1. Catmull, E. & Clark, J. (1978). *Computer Aided Design*, 10 (6), 350-355.
2. Schroder, P., Zorin, D., DeRose, T., Forsey, D. R., Kobbelt, L., Lounsbery, L., & Peters, J. (1998). Subdivision for modeling and animation. *SIGGRAPH 98 Course Notes*.
3. Stam J. (1998). Exact evaluation of Catmull-Clark subdivision surfaces at arbitrary parameter values. *Proc. SIGGRAPH 98*.



Designers of virtual reality (VR) applications need effective tools to develop applications that meet the constraints of virtual environments.⁵ Many current tools have limited flexibility and require in-depth knowledge of the underlying VR system.² We created VR Juggler,¹ an open-source virtual platform for VR application development, to provide a standard environment for application development, testing, and execution. VR Juggler’s design allows for features such as dynamic reconfiguration, input abstraction, and performance monitoring. This sketch describes these features and how they benefit VR developers.

DESIGN APPROACH

VR Juggler utilizes an object-oriented software design that uses standard interfaces to encapsulate VR components based on their functionality, thereby isolating applications from the underlying hardware details. For example, a tracking device uses a positional device interface to hide all hardware details. An application only cares about the type of data it sends and receives; the library handles all the details of accessing the actual hardware.

The component abstraction supports dynamic reconfiguration of applications at runtime. Users are able to add, swap, or modify any component without affecting running applications. This reduces downtime because an application does not have to restart to make changes to a single component, which can be very useful when dealing with unstable hardware or setting up a new VR system.

VR Juggler also provides a framework for application development. It supplies a well-defined interface that a developer realizes to create an application object. This interface assists development by providing a contract that guarantees when each part of the interface is executed and what the system’s state will be at that time. An added benefit of the application interface is that an application is merely another system component. An application object may be added, removed, or exchanged at runtime. VR Juggler also allows multiple applications to run concurrently.

VR Juggler scales across multiple platforms and VR systems. Developers can create applications in a simulator using the resources available at their desktops without requiring access to an actual VR system until the application is near completion.

IMPLEMENTATION

At VR Juggler’s center is a portable microkernel,³ a minimal core that coordinates the entire system (Figure 1). The microkernel builds on an operating system abstraction layer that hides platform-specific primitives such as threads and semaphores behind common interfaces to provide portability.

Most of VR Juggler’s services are provided and encapsulated by components called Managers. They plug into the microkernel and can be added and removed at runtime. Examples are the Input Manager that regulates input devices and the Draw Managers that handle rendering tasks for specific graphics APIs.



Figure 1: Architecture overview.

Applications are components that the microkernel and Draw Manager execute using the standard application interface defined by VR Juggler.

CONCLUSION AND FUTURE WORK

VR Juggler provides a flexible platform for VR application development. Developers benefit from its dynamic reconfiguration capabilities, extensibility, portability, input abstractions, and componentization. VR Juggler is an open-source project and currently runs on many platforms. Current work includes using Bamboo⁴ as a component framework and extensions to support clusters and distributed applications. See www.vrjuggler.org/ for project status.

References

1. www.vrjuggler.org/
2. Bierbaum, A. & Just, C. (1998). Software tools for application development. *ACM SIGGRAPH 98 Course Notes*, 3.1-3.45.
3. Buschmann, F., Meunier, R., Rohnert, H., Sommerlad, P., & Stal, M. (1996). *Pattern-oriented software architecture: A system of patterns*.
4. Watsen, K. & Zyda, M. (1998). Bamboo - Supporting dynamic protocols for virtual environments. *Proc. IMAGE 98*.
5. Cruz-Neira, C. (1995). *Virtual reality based on multiple projection screens: The CAVE and its applications to computational science and engineering*. Doctoral dissertation, University of Illinois at Chicago.

Motion-capture based animation has received increased attention over the last decade, and many new techniques have been proposed for editing and retargeting motion capture data. However, most of these techniques do not specifically address a feature of the data that is the essence of life-like movements: Randomness. Every movement is unique and different every time we execute it. For example, if we execute the same gesture twice, it will not come out exactly the same each time.

We propose a technique that uses motion capture data as input to create synthetic data that keeps the important subtleties of the original data while introducing random variations. The generated data should be statistically similar to the original data and yet not be an exact copy. We use a multi-resolution approach, as it provides a good way to separate out long term fluctuations such as variations in stride length from high frequency noise.

EXPERIMENTS

Our input data is a set of several recordings of the same motion, in this case, several kangaroo hop cycles. Such a database is necessary to provide samples of how the motion naturally fluctuates. The motions of the hip, knee, and ankle joint and the x and y translations are captured using an adaptation of a video-based markerless tracking technique².

We apply a wavelet decomposition to the data³, and in each frequency band use a kernel based representation of the conditional probabilities of various features in the data such as the value of the data at the same time in lower frequency bands, the value at the same time in other

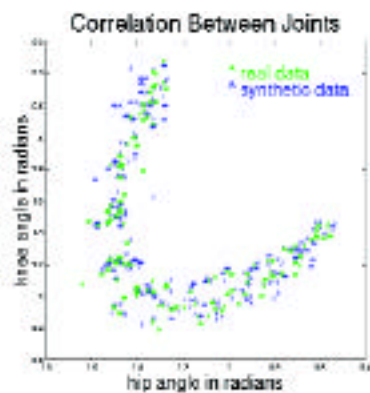
joints, and the value of the data points that precede and follow the one in question. The kernel-based⁴ method is useful because it is a good compromise between the two opposing extremes: a) directly using snippets of motion capture data through a template lookup, which is equivalent to playing one of many “canned” motions; b) representing the data by a low-dimensional parametric density, which could “wash out” important subtleties.

To synthesize the data, we create the lower frequency bands with random numbers scaled to match the real data. The rest of the levels are computed one at a time based on the kernel-based densities. An initial guess is made by computing the expected value of each point in the frequency band given the kernel-based densities for the previous point and lower frequency band. This band is added to the bands previously synthesized to create the synthetic output of a low pass filter, and an iterative optimization technique is applied to find a local maximum in the probability density for this data. For subsequent joints and the translations, the same technique is applied, except that correlation to other joints is also used as a feature. The results are motion curves that drive realistic animations of hopping creatures.

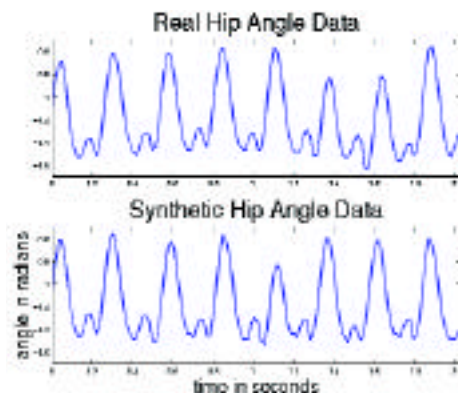
http://graphics.stanford.edu/~pullen/motion_texture

References


1. Motion Texture was originally coined by Ken Perlin
2. C. Bregler, J. Malik, Tracking People with Twists and Exponential Maps, *Proc. Computer Vision and Pattern Recognition*, 1998.
3. G. Strong. *Wavelets and filter banks*. Wellesey.-Cambridge Press 1996.
4. C. Bishop, *Neural networks for pattern recognition*. Oxford, 1995.



It is important that the synthesis technique maintains the correlation between joint angles that are present in the original data. This correlation plot shows that the synthetic data compares well with the real data.



Real and synthetic hip angle data.



Example frames from original data showing the captured joint angles.



Example frames of the output animations

GeoSCAPE is an application that facilitates capturing a 3D visualization of an archeological excavation using HandSCAPE, an orientation-aware digital tape measure (Figure 1). HandSCAPE is an input device for digitizing field measurements to visualize the magnitude and direction of the resulting vectors with computer graphics. Using embedded orientation-sensing hardware, HandSCAPE captures relevant vector information on each linear measurement and transmits this data wirelessly to a remote computer in real time. The GeoSCAPE software analyzes this data, using trigonometry to define vector equations in Cartesian coordinates. Using certain heuristics, these vectors are applied by GeoSCAPE to recreate an archeological excavation rendered in computer graphics.

THE APPLICATION

An archeological excavation is examined and documented primarily using handwritten measurements, notes, and photography, which are often irreplaceable during the on-site excavation. To fulfill the need for efficiency and functionality in archeological excavations, GeoSCAPE visualizes the primitive 3D objects created by measuring. It occurred to us that a new application of HandSCAPE could integrate on-site and laboratory archeological research. By using HandSCAPE to perform measurement tasks with efficiency and speed, the user automatically defines the layout of the excavation on the laptop (Figure 2). As 3D spacing is an important factor in analyzing archeological work, GeoSCAPE provides an important step toward linking the physical archeological excavation with a virtual rendering on-site.

3D VISUALIZATION

Our current implementation of GeoSCAPE defines the relationship between the physical space and the virtual space as follows. The “anchor point” in physical space is one corner of the excavation and is represented in the application as a white sphere. This is also the origin of the 3D coordinate system in the virtual space. From this anchor point, a new reference frame is defined by making a vector in physical space and sending the data to the computer by pressing a button on the HandSCAPE device. Within this reference frame, three additional measurements are taken. These are used to define a cubic figure in the

virtual space that represents the location of an artifact in the physical excavation. Issues of denoting perspective were also important. Measurement grids were added, as was a 3D compass. The compass, seen in the upper right corner of the screen shot, has its 3D rotation tied to the user’s movement of the virtual excavation. Other functionality introduced here includes a map of the excavated region, display of coordinates and dimensions for the artifact bounding boxes, the ability to create a multi-tier excavation site, and full 3D navigation.

FUTURE EXTENSIONS

To provide a single source of access to information regarding the excavation, we are currently working on connecting GeoSCAPE to an archeological database with notes, photographs, and other relevant information. This database is used during the excavation to aid in reconstruction of the site by referencing stored information. This would allow for immediate access of information pertinent to the objects being viewed and allow the user to better visualize the site. Another idea being pursued is integration of a GPS positioning device into the HandSCAPE tool. GeoSCAPE could link this to a map program to change the map view automatically depending on location and store the data in the database.

ACKNOWLEDGEMENT

We would like to thank Louise Krasniewicz, Richard M. Leventhal, and Ken Stuart of the Cotsen Institute of Archaeology at UCLA for the discussion on the archeological application. We also thank the members of the Tangible Media Group and our colleagues of the Things That Think consortium at the MITMedia Lab for their support of this research.

References

1. Dillon, B.D. (Eds). (1993). *Practical archaeology: Field and laboratory techniques and archaeological logistics*. Los Angeles: University of California Press, 33-38.
2. Lee, J., Su, V., Ren, S., & Ishii, H. (2000). HandSCAPE: A vectorizing tape measure for on-site measuring applications. In *Proc. of Conference on Human Factors in Computing Systems CHI’00*, 137-144.



Figure 1: The HandSCAPE tool (12 x 12 x 4.5 cm).




Figure 2: New Approach: GeoSCAPE is an on-site archeological application of HandSCAPE, a digital tape measure with custom sensing electronics. It communicates through a RF signal and performs graphics rendering.

Conventional human interface devices handle only 2D information. Our purpose is to build a 3D interface with a depth-image flow in real-time. A new kinetic and 3D image input device motion processor¹ can detect an object's distance from the camera to the object in real-time, its 3D posture using image moments, and normal vectors of the object's surface captured by the motion processor. It can also check the detection error of posture angle within five degrees. We apply this framework to 3D posture detection of hand gestures for a virtual ping-pong game. This application shows that a 3D interface based on real-time gesture flow is essential to natural and intuitive communication between humans and computers.

Current posture input devices for 3D interfaces using hand gestures². have some problems. The motion-capture system requires users to wear cumbersome sensors. A vision-based system using multi-CCD cameras is considerably affected by background conditions when it extracts users' hands. A depth-image capturing system using a range finder takes a fairly long time to capture a frame. We use a motion processor (Figure 1) as the image input device, to illuminate the hand with infrared light, and the reflected light is captured by a custom image sensor (Figure 2). We can extract the user's hands from the background and capture the hand's depth image in real-time (50 fps).

To construct the object's 3D model from the captured depth-image, we represent each image pixel $I(i,j)$ as a 3D point $(i, j, I(i,j))$ and each object surface as a group of quadrangle patches. Each patch has four neighboring 3D points (Figure 3). The 3D normal vector of the hand surface ($N_{surface}$) is calculated by averaging each normal vector of

quadrangle patches that consists of object surface. Assuming that the hand surface is flat ("open"), the hand's gradient posture toward depth direction is calculated as a rotational displacement angle from the frontward posture, where the normal vector is the front direction $N_r=(0, 0, 1)$. This means that the hand's rotational angles along the X/Y axis are calculated to transform the frontward normal vector N_r to the hand surface's normal vector $N_{surface}$. We can detect the flat object's rotational posture angle along the Y axis within five degrees of error.

The result, our ping-pong interactive game application, demonstrates our gesture-based 3D interface. The virtual paddle's screen position and 3D posture reflect the detected user's hand position and 3D posture in real-time (Figure 4). Users can manipulate the virtual paddle's position and posture to intuitively control ball behavior by changing their own hand surface as they would do with a real paddle in real ping-pong game.

The 3D interface described here can also be applied to 3D designing or modeling application areas as an intuitive manipulation interface.

References

1. Numazaki S. et al. (1998). A kinetic and 3D image input device. CHI '98, 237-238.
2. Pavlovic, V.I., Sharma, R., & Huang, T.S. (1997). Visual interpretation of hand gestures for human-computer interaction: A review. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 19 (7), 677-695.



Figure 1: Outlook of Motion Processor




Figure 2: Principle of Motion Processor




Figure 3: 3D representation of object surface




Figure 4: Screen image of "Ping-Pong" game

“Shrek,” PDI/Dreamworks’ second CG-animated feature film, represented new challenges in fluid simulation beyond what was achieved for “Antz.” New methods had to be developed in order to simulate heavier, thicker fluids (such as mud), improve the interaction with the characters, and mix two liquids of different properties in the same simulation.

Building on top of the fluid simulation implemented for “Antz,” a number of systems were created to help solve a variety of problematic situations. Problems were broken down into three stages: simulation, wet-map generation, and rendering.

The simulator system allows an animator to use PDI’s proprietary fluid simulator tools (written by Nick Foster) to lay out initial fluid conditions in a shot and determine which characters and objects will be obstacles. It manages oversight of big environments and optimizes those environments by letting the animator run several smaller simulations only where they are needed. The animator can also produce other motion using PDI’s in-house scripting language or other tools to augment the fluid simulator data. The simulator system keeps track of these and handles merging them together for rendering.

The fluid tools were also improved to manage a wider variety of situations. They now allow an animator to mix fluids of different properties in the same simulation. The other important area of improvement was to allow moving characters to interact with the fluid in a realistic and fairly automatic manner without increasing the resolution of the simulation.

Used to generate maps on objects to indicate where they have been wet, the wet-map generator takes the particle output of the simulator and detects the particles that collide with the object. A copy of those particles is attached to the character for the remainder of the animation. These particles are then used inside a shader at render time to indicate wet regions on the object. This allows objects to easily go from dry to wet in the same shot.

The rendering system uses the output of the simulator and prepares it for rendering. The output of the simulator is a set of particles that can be textured, lit, and rendered using PDI’s proprietary particle renderer (initially developed for “Antz”). For “Shrek,” the renderer was extended to use the same lights as the normal renderer to ease integration of particles into the scene. This renderer is extremely fast, making it easy to render over a million particles in just a few minutes.

The rendering system is also able to generate a surface from the particle set. This is crucial when doing fluids, such as water, which don’t have a particle look, especially when they are flowing slowly. The surface can have different characteristics depending on the type of fluid it represents. This allows one surface to be generated from a dataset of multiple types of fluids with the right geometric properties in the correct locations. This system also assigns different material properties to different regions of the surface.

CONCEPT

Spatial architecture of information on the Internet is visualized in GINGA (Global Information Network as Genomorphic Architecture), where metamorphic relationships in cyberspace are manifested in the form of an "Information-scape." GINGA is proposed as another browsing system to explore the World Wide Web more spatially and intuitively. Web resources suggested by Internet users are reconfigured with cyberspace codes into the Worlds of GINGA. Users can access these Worlds through a VRML interface and search information by each world's classification system and geometrical features.

SYSTEM

GINGA has a server/client architecture, and users access the GINGA server through Web browsers. A user's individual system consists of three major functions: avatars, editor, and personal archives (see Figure 1). Avatars can cruise and search the worlds and communicate with other avatars. The editor is a user's interface for selecting worlds and managing avatars as "cyber-librarians."

GINGA is expected to work as a multiple visual search engine.

NINE WORLDS

GINGA's information-scape appears as any of nine main worlds:

1. *Nebula*: Information is distributed in 3D space according to IP addresses of host computers and Web servers. These formations reflect topological relationships in the World Wide Web. The Nebula represents the default condition of GINGA (see Figure 2).
2. *Ring*: Information is assembled on the rotating rings. The diameter of each ring is determined by the frequency of that information's update and its width by the volume of the information.

3. *Network*: Information linkage is emphasized in the Network. Network nodes serve as sites of intersection for multiple subjects. The size of the connecting paths reflects the number of transmitted bits.
4. *Forest*: Groups of a tree directory system are represented in the forest. The height of each tree corresponds to the number of levels of hierarchy from the subject to the broadest category.
5. *Strata*: Information accumulates in chronological order in the strata. Each fibrous segment within a stratum contains strata of the next smaller period of time.
6. *Text*: The text is an exclusive archive for text information. As Jorge Luis Borges noted in "The Library of Babel," it contains an infinite number of infinitely thin leaves of text.
7. *Image*: The image is a world of pictures, maps, photos, movies, and other visual information. The floating image is the first of a group of related images that can be accessed as you enter Image mode.
8. *Sound*: Digital sound data are transformed into notes of visible bits. These bits spiral around a vast number of light columns that correspond to a type of music, musician, or musical instrument.
9. *Cemetery*: The cemetery is a cyber-necropolis for unused or dead avatars. When users die, their avatars are usually buried here. Memories of dead avatars can be accessed, if necessary.

CYBERSPACE CODES

Each GINGA World can be visualized and configured by cyberspace codes (see Figure 4).

www.plannet-arch.com/ginga.htm




Figure 1: System of GINGA




Figure 2: Visualization of IP address



Figure 3: User's interface of GINGA

allocation	data identifier such as IPs
network	links between websites
depth	cataloging specificity
detail	file, summary, entire text, etc.
cycle	update cycle
date	updated or last modified date
type	text, image, movie, sound, etc.
size	size of information
velocity	accessibility, speed of connectivity

Figure 4: Cyberspace codes




Figure 5: Nine worlds of GINGA

In September 1999, ILM was approached by Gatorade to produce a new commercial. We accepted the responsibility of making a five-shot test in order to demonstrate that we could not only make a believable raptor, but also a vicious basketball-playing one! After the test, the job was awarded, and we received the boards for the 19 shots of the spot.

The director, Steve Beck, wanted Vince Carter to play against another player (raptor stand-in) in all the shots. This created a more realistic and dynamic performance but also brought the challenge of removing the stand-in player and replacing it with the raptor. CGC's (Computer Graphics Commercials) painters were responsible for either completely removing the stand-in or blending the stand-in's parts into the background plates.

On set, VFX supervisor Pablo Helman captured all the references needed, plus a library of liquid elements. For reference purposes, we used a raptor's head sculpted by Randy Dutra, who was also in charge of animation direction (with Dan Taylor as lead animator). Leandro Estebecorena was appointed lead technical director for the show, and he was also responsible for making a setup to solve all the technical problems (from lighting to liquid simulation) for the 19 shots of the spot.

To final the model on time, we used a shader to evaluate normals, allowing faster detection of creases and artifacts on the model. New muscle development and control were needed, and a library of corrective shapes was created for each shot.

The raptor was supposed to sweat (drops would be emitted when the raptor hits Vince), and liquid would drip from its mouth when it drank Gatorade. We used practical water elements for some shots (composed with ILM custom tools), and generated “CG sweat and liquid” for those shots in which close water-raptor interaction was required. We divided the “CG sweat and liquid” problem into three types: Sweat emission, dribbling sweat, and liquid simulation. Nigel Sumner, who was responsible for the R&D of liquid simulation for the show, used a combination of custom scripts and Maya to solve the first two and ILM custom applications to create the liquid simulation.

The CGC technical directors faced, and solved, a lot of problems to seamlessly integrate the Raptor in the game with Vince: color-corrected and sharpened plates with extremely contrasted grain, flashy lighting, players coming toward the camera through a foggy environment, flare occlusion, volumetric lighting and shadows, creation and integration of liquid, etc.

Scripts were developed to generate proofsheets for color continuity evaluation between all the spot's shots on a daily basis. Additional problems (like adding blue drops of sweat on Vince or creating additional mattes) were solved using Sabre, an ILM high-speed compositing system that incorporates Discreet's Inferno-Flame software.

The spot started production in December 1999 and finalized on 4 February 2000, just in time for the NBA All-Star game, Sunday, 13 February 2000.

A hardware-accelerated algorithm is presented to compute the error between a silhouette in an image and the silhouette of a projected 3D model to estimate the camera position.


Since rendering algorithms have become more sophisticated in recent years the demand for realistic models has increased, especially for models of real world objects. As a result, not only the 3D shape but also the surface color or reflectance is of major interest. The 3D model is commonly obtained by a scanning device producing, for example, a triangle mesh. Color information is frequently collected during a second process using a camera. Generating the entire model then requires combining the 3D shape and the 2D images/textures. The camera position and orientation relative to the 3D mesh must be computed for each image in order to stitch the textures onto the 3D surface. In previous work,^{1, 2} the camera settings are determined by minimizing the error between the outline of the object in the 2D image and the outline of the projected 3D model.

In contrast to previous algorithms, where the error is estimated by a sum of pixel-to-outline distances, our approach directly measures the area enclosed by the two outlines. Since the projection of the 3D model

can be computed using graphics hardware, it is very easy and efficient to obtain this area measurement. As a preprocessing step, the 2D image is segmented into pixels that belong to the background (black) and those covered by the object (white). For each camera position, at first the 3D model is rendered white in front of a black background, and the result is combined with the segmented image using a per pixel XOR operation. As a result, exactly the pixels between the outlines remain white, and their number can be counted by simply evaluating a histogram. Since all these steps can be executed in hardware, the non-linear optimization for the camera parameters is significantly accelerated. Further, the optimization is done hierarchically on different image resolutions to decrease the rendering effort. Given the focal length of the camera and a rough estimate for the object distance, it is even possible to obtain the camera position and rotation in a few minutes, without any user interaction.

References

1. Matsushita, K. & Kaneko, T. (1999). Efficient and handy texture mapping on 3D surfaces. In *Computer Graphics Forum*, 18, 349-357.
2. Neugebauer, P.J. & Klein, K. (1999). Texturing 3D models of real world objects from multiple unregistered photographic views. In *Computer Graphics Forum*, 18, 245-256.



Measuring the difference between the photo and one view of the model by the area occupied by the XOR-ed foreground pixels.

For animated subdivision surfaces, where the initial mesh vertices can move freely, all vertices of the final mesh are usually recomputed each frame. We present a simple heuristic technique that not only culls backfaces, but also avoids computing a large number of vertices of these backfaces. The basic idea is illustrated in 2D in Figure 1, where we consider the backfacing line segment A. Since both its neighbors (B and C) are backfacing, we can prove that A' and A'' also will be backfacing, and thus we avoid computing them. C'' and B' are also provably backfacing, but are needed to properly subdivide their neighbors C' and B'', which may become frontfacing.

In 3D, we concentrate on Loop's scheme.^{2,3} The neighbors of a triangle Δ are those triangles that share either an edge or a vertex with Δ . Our heuristic is: if a triangle Δ and all its neighbors are backfacing, then all new subdivided triangles "inside" Δ are also assumed to be backfacing. We call Δ strongly backfacing, and a triangle with all neighbors being strongly backfacing is called totally backfacing.

In each subdivision step, we determine the status (front, back, strong, totally back) of all triangles that have not yet been culled. Totally backfacing triangles are excluded from all further computations. This conservative culling is necessary because part of the invisible neighborhood of the front side is needed (strongly backfacing triangles must be subdivided because the first "row" of backfacing triangles may need their neighbors).

To accelerate further, we determine frontfaces similarly. The new subdivided triangles of a completely frontfacing triangle (whose neighbors are all frontfacing) are all completely frontfacing, and the status computation is omitted. Our algorithm is shown in Figure 2.

Mathematically, there are no guarantees that meshes with grossly varying facet sizes cannot produce "flipped" facets with orientations almost opposite to that of their neighbors. Such sign reversals, if not present in the initial control mesh, are clearly unwanted artifacts that should be suppressed, and that is exactly what our heuristic does.

Results: for a high enough degree of subdivision (for example, when each resulting primitive must be smaller than half a pixel,¹ we have observed reductions in rendered polygon count close to the theoretical limit of 50% and speed-ups of close to 100%. The overhead for status bookkeeping is one byte per triangle, which is typically negligible, and well worth it, since we now store only 50% of the triangles.

Code overhead is 5-10% of the execution time.

References

1. DeRose, T., Kass, M. & Truong, T. (1998). Subdivision surfaces in character animation. *Computer Graphics (Proc. SIGGRAPH 98)*, 85-94.
2. Hoppe, H., DeRose, T., Duchamp, T., Halstead, M., Jin, H., McDonald, J., Schweitzer, J. & Stuetzle, W. (1994). Piecewise smooth surface reconstruction. *Computer Graphics (Proc. SIGGRAPH 94)*, 295-302.
3. Loop, C. *Smooth subdivision based on triangles*. Masters Thesis, Department of Mathematics, University of Utah, August 1987.




Figure 1: 2D example.




Figure 2: The camera that does backface culling is located in the upper right back. a) Initial worm mesh (color codes: white=totally back, blue=strong, green=back, red=front, yellow=completely front). b) After one subdivision step: the white triangles from the previous step have been discarded. c) Two subdivision steps. d) Three subdivision steps (what our algorithm sends to the graphics pipeline). e) What normally is sent to the graphics pipeline without our technique.

i2TV (interactive Internet TV) is an experiment in researching new dramaturgical formats, media integration, and models for medial staging of online/on-site events, such as integration of Internet participants in an event taking place at a real physical location. The challenge is to create a setup in which online and on-site participants are equal partners, while retaining the specificity of both situations (online, on-site). The different space-time conditions of online and on-site participants pose the biggest problem: no less than an emerging architecture of time/space is questioned. To achieve this goal, the experiment concentrates on the combination of new dramaturgical concepts in an architecture of medial elements. Exploring and developing such concepts relates to lessons and experiences from different fields of the arts, from traditional theater and 1960s “participatory performances” to interactive situative environments and mixed-reality stages.

This sketch discusses issues of presence and presentation underlying the i2TV system and its technical realisation, as well as experiences in public rehearsals: the Memoria Futura symposium at GMD, and the distributed poetry play experiment at the European Media Art Festival in Osnabrück. We discuss open questions and possible strategies for further exploring i2TV as the basis for a medial staging that integrates Internet and on-site participants. The presentation itself tries to demonstrate a new format by using i2TV as a front end for artistic live production.

The i2TV event at the Memoria Futura Symposium focuses on integrating online/on-site participation as an extension of the common conference format and moving it toward new models for connecting people or communities. The on-site discussion integrates a group of invited experts as Internet participants through the i2TV system. It is the proof of concept for i2TV as a real-time conference and archive system, and editorial online environment. The system builds on top of publicly available MOO-WWW technology (enCore Xpress)¹ combined with live audio-video, virtual environments, and digital TV.

The EMAF2000 (European Media Art Festival) i2TV event concentrates on new dramaturgical concepts for online/on-site events. It explores a theatrical concept as a basis for integrating online and on-site participation with the staging in physical space. The distributed poetry play experiment aims to involve the participants in an open workshop and/or a performance situation. Questions addressed are: How can the conflict of “real-time delay” and the conflict of “immediacy of involvement — distance for reflection,” as intrinsic qualities of the online/on-site situation, provide fertile ground for new artistic concepts? How do we create a spatial situation in which both online and on-site participants feel present and involved, even if not in the same way? How do we create a situation that is both theatrically staged and enables reflective discussion, at the same time? How can such an approach help us to open the “black box” of fast-developing technology to a public consciousness?

1. Developed by Jan Rune Holmevik, University of Bergen, and Cynthia Haynes, University of Texas at Dallas, as a platform for academic virtual communities (lingua.utdallas.edu).



Figure 1: Current i2TV Web interface.




Figure 2: Conceptual sketch for a distributed poetry play.

We have been building software to create contour maps from undersea echo soundings for the “Bridge and Barrier” research program of the Museum of New Zealand Archaeozoology Laboratory.

Drawing contours is straightforward once you have a smooth function. The problem is to create that function from discrete data.¹ We use the method of smoothing the data with a filter,^{2,3} but there are millions of depth figures, and the data are very unevenly distributed. When you have a lot of samples close together, their combined weight is larger than it would be if you had only one sample in the region, and the effect is to bias the average.^{4,5}

A finite filter with short range gives results like Figure 1. The light blue area represents a region where we had no data. The client was not happy. Could we use the coastline data to make a guess at the contours where there were no measurements? It seems wrong to invent data in a scientific application, but we perceived that in a sense the client was right. Our “contour” around the edge of the data is artificial, and depths a little outside this area were still likely to be close to the nearest measurements. But a longer-range filter would surely introduce more bias.

Our solution is to resample our smoothed data with a regular grid of samples and then smooth the resampled data with a wider filter. This second level of smoothing creates “reasonable” data in empty areas but also affects the accuracy of the contours where we have data. We can, however, combine the two smoothed functions with different weights to create a smoothed function that follows the data accurately where it exists and follows the long-range “guess” where it doesn’t.

We use a layered structure where each layer carries data at a different resolution. Every entry in this structure consists of a “weight” and a “depth.” In each node of the data structure, we store the sum of depth times weight and also the sum of weights.

We draw our contours from a continuous depth function, calculated at sample points, from the layers. From each layer, we interpolate

a value from the four nearest grid points to our sample point. We combine these interpolated values as a weighted sum where the lower layers have a higher weight than the upper layers. If a region contains no local data, it will acquire an average value from some higher layer but with a low weight. This process is basically the same as described in Gortler.⁶ Our implementation uses a continuous cubic function rather than the “hat” curve.

Storing the weights gives us a real bonus. When we create the weighted sum for the sample points, the sum of weights gives us a measure of how much local data contributes to the value. In Figure 2, we have used this weight data to color the contour lines. Red indicates a high certainty from good local data. Yellow means we have a guess based on coastline and global average depth. Intermediate colors estimate the confidence we have in the contour at that location.

Not all the contours in the originally “good” area are red. Our improved method of plotting has actually revealed areas of uncertainty that we were not originally aware of.

Details are available at:

www.cs.otago.ac.nz/gpxpriv/public_html/Research/contours.html

References

1. Hardy, R.L., (1971). Multi quadratic equations of topography and other irregular surfaces. *Journal of Geophysical Research*, 76 (8), 1905-1915.
2. Perlin, K. (1985). An image synthesizer. *Computer Graphics (SIGGRAPH 85 Conference Proceedings)*, 19 (3), 287-296.
3. Lewis, J.P. (1989). Algorithms for solid noise synthesis. *Computer Graphics (SIGGRAPH 89 Conference Proceedings)*, 23 (3), 263-270.
4. Mitchell, D.P. (1987). Generating antialiased images at low sampling densities. *Computer Graphics (SIGGRAPH 87 Conference Proceedings)*, 21 (4), 65-72.
5. Kirk, D. & Arvo, J. (1991). Unbiased sampling techniques for image synthesis. *Computer Graphics (SIGGRAPH 91 Conference Proceedings)*, 25 (4), 153-156.
6. Gortler, S.J., Grzeszczuk, R., Szaliski, R., & Cohen, M.F. (1996). The lumigraph. *Computer Graphics (SIGGRAPH 96 Conference Proceedings)*, 3-54.

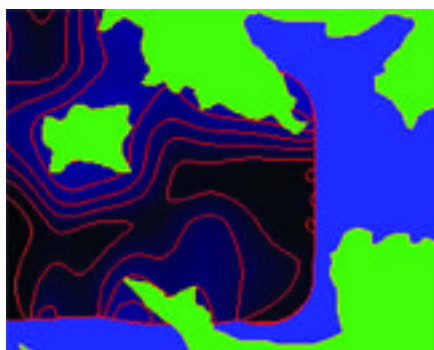


Figure 1: Original presentation ignoring coastline.




Figure 2: Showing uncertainty.

The Direct Rendering Infrastructure (DRI) is an open-source framework that allows applications to directly access graphics hardware while cooperating with the X Window System. The goal of the DRI is to provide high-performance graphics while maintaining the security and stability of the system.

The first application of the DRI has been to implement the OpenGL API in Linux. Linux runs on a wide variety of hardware, including commodity Intel-based PCs, and already has a high-performance implementation of the X Window System created by The XFree86 Project, Inc. Commodity hardware vendors, Linux system developers, and XFree86 developers are all adding features and improving performance at an incredible rate. These systems will make excellent, low-cost, high-performance 3D workstations available in the very near future.

GOALS

The overall goal of the DRI is to provide an open-source solution that allows a system to perform at the full potential of the hardware. Some key aspects of this project include: the benefits of open-source development; support for commodity hardware; high performance; flexibility to support a wide variety of hardware designs; and maximum software reuse, security, portability, and adherence to standards.

DIRECT RENDERING

Direct rendering achieves its high performance by allowing the application to communicate directly, yet safely, with the graphics hardware, in contrast to indirect rendering, where all the graphics commands are sent to the X server, which is responsible for the centralized control of the hardware, to be processed. Direct rendering lowers the latency and increases the bandwidth of the communication between the application and the hardware.

DRI IMPLEMENTATION

There are three components to the DRI implementation. The first is a 2D driver within the X server. It performs any initialization required and communicates the state of the window system to clients as it is modified. The second is a 3D driver, which is part of a shared library linked directly to the client binary. It communicates directly with the hardware to render primitives on behalf of the application after authenticating itself to gain access to restricted resources and requesting configuration information from the X server. The third is a kernel-device driver. The device driver provides the secure authentication mechanism, synchronizes access to the hardware, and provides an array of services that the 3D driver could use, if necessary, to access the hardware (for example, virtualization of the hardware, DMA dispatch, and context switching).

DRI SERVICES

The first service the DRI provides is coordination among the processes that access the hardware. The DRI provides three mechanisms for this communication: a shared memory area for fast access to data shared by all three components of the DRI; an X protocol extension, which is slower, but allows for large amounts of data to be requested by the 3D driver from the X server, and a kernel-device driver service that arbitrates between multiple processes.

The second basic service provided by the DRI is direct access to the hardware. The interface the kernel driver provides is highly dependent on the hardware implementation itself. In some cases, direct access to portions of the hardware may be a security violation. In these cases, the device driver may provide its own interface to the hardware that can perform appropriate security checks before performing an operation.

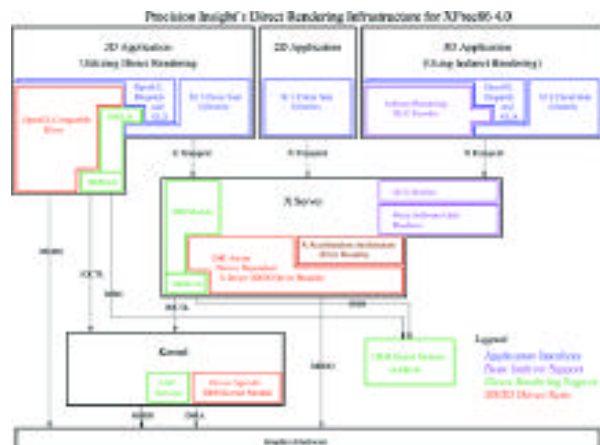
The third service provided by the DRI is authentication. When a client wants access to the hardware, it must first connect to the X server, which provides a number of authentication mechanisms that allow fine-grained access. Only those clients that can connect to the X server are given direct access to the graphics hardware.

FUTURE PROJECTS

Precision Insight, Inc. has created several open-source implementations of DRI clients, but the DRI is a work in progress, and there are many features that we would like to include in future versions, including multi-head support, sophisticated texture and off-screen memory management, support for 3D hardware that includes transform and lighting capabilities, and other features. The DRI can also be used for applications other than OpenGL. The next implementation will include support for video applications interacting directly with the hardware.

CONCLUSIONS

The DRI is a general-purpose framework for directly accessing graphics hardware. The generality of the solution is proven by the diversity of the hardware implementations created so far and by the support of a variety of graphics vendors in the industry. The implementation is already quite functional and will be improved as more hardware is supported and as features are added and improved. The value of this open-source solution can be seen by the speed in which it was developed and adopted by the Linux community.



In Paul Verhoeven’s “Hollow Man,” a modern-day twist on the invisible man, Sony Pictures Imageworks was tasked with creating a CG version of the actor, Kevin Bacon. However, because of the way our title character transitions between visibility and invisibility, this required replicating not only the exterior skin, but also all the underlying anatomical structures, getting them to move in a realistic, believable, manner as well as efficiently managing the extremely large dataset that constituted our synthespian.

“Physiquing” is a term used at Imageworks to describe the deformation of the various anatomical body parts in a physiological way. Global issues had to be addressed, such as volume preservation and interpenetration. In addition, each organic type had its own specific set of problems. For example, for muscles, we had to ensure that origin and insertion points stayed attached to the bone and pay attention to how the tendons reacted compared to the muscle body.

While we did strive for anatomical accuracy, it was purely in an aesthetic sense. Animator-required interactivity, the production schedule, and the aesthetic flexibility required in making a motion picture prohibited us from simulating the proper mechanical nature of a living human being. For our digital Kevin, the movement of the skeletal structure drove the deformation of all the geometry. In real life, deforming muscles move the supporting skeleton.

The primary animation and deformation framework was built around Alias|Wavefront’s Maya. Using the open architecture of Maya, custom plug-ins were written for skin deformation and relaxation, stitching, collision detection, and proprietary deformers. Extensive use of the MEL

scripting language allowed us to extend many of the baseline character tools, which allowed us to constantly rebuild the digital actor and customize a setup for a particular shot if the need arose.

Deformation of the internal structure was achieved through a layered set of deformers. The primary set was procedurally driven based on joint rotations. Additional deformers and associated controls were layered on top of this procedurally driven set to allow for finer deformations, corrections, and changes requested by the animation supervisor, visual effects supervisor, or director. Animators had the ability to implement details such as isometric contractions and specific volume changes to emphasize a particular action.

Facial muscles, however, were deformed based on deformation of the skin, which was achieved primarily through blend shapes. Points on the face were analyzed, and their movement drove the set of facial-muscle deformers. This allowed an animator to use a single, interactive interface to create facial skin and muscle animation. As with the body, additional deformation controls were layered on top of the base set to allow for finessing a shot as needed.

Simulating the skin and interconnective tissue required custom deformers. These Maya plug-ins allowed the skin to move based on the underlying musculature and allowed it to relax and slide over that surface. It also allowed seams of different resolutions to stitch together.

By completing nearly 100 shots with our digital human, the system has shown it is capable of replicating a living being, which opens the door for future synthespians.

Certain types of character elements should move in a complex or dynamic way. Hand-animating these elements provides explicit control but requires a large amount of animation time. Simulating these elements can be cost-effective, but controlling a dynamic system can be complicated. PDI is integrating dynamics into a character, utilizing traditional animation controls as the simulation's main input, with enhanced predictability and productivity. This produces believable dynamic motion but allows for more explicit and traditional controls in cases where the animator or director desires to see something more specific.

Elements that work best in this kind of system are parts of a character's skeleton that can be described as a series of joints attached at one end, each joint having three degrees of freedom in rotation. Classic examples include hair clumps, tails, antennae, and long ears. PDI's general dynamics system, "strand," can be integrated into these portions of a character as an animation tool that provides complex motion, increasing productivity while using traditional animation controls.

The main input to the strand system is the animation of the character itself. In the case of a tail, the animator can concentrate efforts on other parts of the character, allowing the dynamics to solely determine the tail's motion. If the animator desires a more specific tail motion, an arbitrary number of tail poses can be set through traditional animation controls to help direct the simulation. The animator is able to further tune the motion through a few crucial dynamics controls like gravity, pose influence, and kinematic-to-dynamic blending.

Pose influence defines how closely the dynamic motion adheres to the kinematic, or animator-defined, motion. The dynamics can be constrained to stay close to the posed strand by increasing its influence, or allowed to have a looser dynamic quality by "dialing down" the influence.

Although the simulation is calculated for the duration of a shot, a blend control provides flexibility to switch between kinematic and dynamic motion. This control solves several problems. The dynamic solution is driven by the kinematic motion, blending from kinematic to dynamic blends between similar poses, so these transitions can be easily concealed.

In summary, PDI developed and integrated a pose-driven dynamics system as part of character setup. Portable between characters and productions, the system uses traditional animation controls as the main input, making it an intuitive tool for animators, enabling them to quickly create complex and predictable motion for many character elements, which dramatically increases productivity.

An implementation of the OpenGL 2D/3D graphics API was developed for use on the Palm Computing platform, which allows easy porting of graphics code from the desktop onto the smaller, mobile personal digital assistant (PDA) platform. The project was challenging in many respects due to hardware limitations and the small screen size of the PDA, and it provided an interesting set of new problems to overcome in the area of interactive 3D graphics display. This sketch illustrates the challenges overcome in porting OpenGL to the Palm platform, demonstrates the usefulness of 3D graphics on the Palm, and shares the library code base with the 3D graphics development community through the LGPL license under which the developed library was released.

The work began with inspection of the OpenGL headers. The goal was to reverse engineer the function calls and enumerations of the OpenGL library. It was known that the Palm OS API includes a simple line-drawing function, so the implementation had to handle all conventional 3D graphics functions, including geometry transforms, perspective transforms, screen coordinate transforms, illumination, shading, texture mapping, surface material properties, color models, and anti-aliasing. Further research revealed Palm OS support for 4-bit grey-scale colors, so the limits of the display were 16 color at a screen resolution of 160 x 160. An additional hardware limitation that came into play in latter stages of the project was the fact that the Palm's Motorola processor generally runs at around 20 MHz, so not only was the physical size of the display constraining, but the computation speed was also a factor in designing

an efficient implementation. Given these hardware constraints, the process of implementing the OpenGL function calls was constantly governed by performance/quality tradeoffs that are typically not a concern in today's world of fast desktop machines and hardware rendering cards. The overarching goal of the project was to allow existing OpenGL code to compile transparently in Palm development environments with similar screen-display representations, so that the act of porting was confined to only interface-specific functionality (that is, by design, not handled by the OpenGL API anyway).

This project successfully used the base pixel drawing functions of the native OS, combined with an implementation of standard 2D/3D graphics algorithms, to provide a set of OpenGL functions for Palm programs. The geometric and perspective transforms are treated as 4x4 matrices as in the OpenGL model, with translation, scaling, and rotation supported. The supported primitives are the OpenGL set (GL_POLYGON, GL_TRIANGLE_STRIP, etc.), and the lighting model includes up to eight lights with ambient, diffuse, and specular terms. The color model is limited to 16-color grey-scale, but the colors are stored in RGB format for future expansion to the 256-color Palm unit and mapped to the closest grey in the 16-color palette. Texture mapping is accomplished by loading in bitmap images from a database held on the Palm unit and performing per-pixel texturing in the polygon flood-fill stage. Both Gouraud and Phong shading were implemented, and either can be selected upon compilation of the library.

Recently, graphics hardware has increased in capability and is available even on standard PCs. These facts have encouraged researchers to develop hardware-accelerated methods for rendering realistic images. One of the important elements in enhancing reality is the effect of atmospheric scattering.¹ Examples of such effects include light beams from spotlights and shafts of light from the sun's rays. The purpose of this sketch is to develop a method for displaying shafts of light at interactive rates by making use of graphics hardware.

BASIC IDEA OF RENDERING SHAFTS OF LIGHT

Our method treats the effects of scattering due to point light sources and parallel light sources. The intensity of light scattered from particles has to be integrated along the viewing ray to display the shafts of light. Figure 1 shows the basic idea of our method. Virtual planes are placed in front of the viewpoint for the integration. Each virtual plane is parallel to the screen and represented by a lattice mesh (see Figure 1). First, the intensity of the light scattered at each lattice point is calculated. The total intensity is the sum of the intensities of all the virtual planes. It is computed by rendering the virtual planes with an additive blending function. The luminous intensity distribution of the light source can be taken into account by mapping a light map texture² onto the virtual planes with a multiplicative mapping function (see Figure 1).

SHADOWS IN ATMOSPHERE

Shadows cast on particles in the atmosphere are very important; when there are objects in an illuminated volume, non-illuminated parts arise within it. To display the shadow in the atmosphere, non-illuminated parts of the virtual planes must be detected. To achieve this, we have implemented the idea of the shadow map³ by creating shadow textures for each virtual plane. That is, before rendering each virtual plane, an image of the shadow cast on it is created. The image is then used to mask the non-illuminated parts of the virtual plane by mapping it as a texture with a multiplicative blending function.

RESULTS

Figure 2 shows an example of a teapot illuminated by three spotlights. We see shadow effects due to the teapot; the shape of the shaft of light is clipped. Figure 3 shows another example. Mountains shut out parts of the sun's rays, causing shafts of light along the silhouettes of the mountains. Frame rates for these examples are several images per second on the average.

References

1. Nishita, T., Miyawaki, Y., & Nakamae, E. (1987). A shading model for atmospheric scattering considering luminous intensity distribution of light sources. *Computer Graphics*, 21 (4), 303-310.
2. Blythe, D. (1999). Advanced graphics programming techniques using OpenGL. *Course Note 29*, SIGGRAPH 99.
3. Segal, M., Korobkin, C., Widenfelt, R. V., Foran, J., & Haerberli, P. E. (1992). Fast shadows and lighting effects using texture mapping. *Computer Graphics*, 26 (2), 249-252.




Figure 1: Basic idea of rendering shafts of light.




Figure 2: Light beams from spotlights.



Figure 3: Shafts of light from the sun's rays.

Dick Higgins, the late Fluxus artist, once discussed a simple problem in hermeneutics, the 17th century word used to describe methodologies for interpreting the Bible. After several centuries of using vernacular texts, theologians and historians became curious about the authors' original intentions. This study would help them interpret the original meaning of words and phrases. Higgins compared this idea with the metaphor of horizons to describe the relationship of a viewer's horizon to that of an art object and the horizon of the artist. These horizons are completely different, but for a brief moment when looking at an art object, horizons come into contact and fuse before separating again. Higgins believed that if after leaving a work of art, it has shifted your horizon a little, then the work is somehow meaningful.¹


Language can similarly be considered a tool where information is, to varying degrees of adequacy, able to represent ideas from sender to receiver. In this context, language creates a sort of agreed negotiation of communication. However, what happens when the physicality of these negotiations change? What happens when a photograph of our grandfather is no longer a fixed paper object but a manipulatable digital image on a CD? A comparison could be made an earlier transition, when the ancient scroll replaced the portable book/codex form as a better disseminator of information, where form successfully followed function. Today, designers are claiming that form follows imagination.² What happens when opening a book turns on a movie, not because it's functional, but because our negotiations are no longer fixed in form following function. It is here that Higgins' multiple horizons are useful when working with interactivity. In the case of mediated tools, multiple horizons ultimately fuse the artist, the interactive device, and the viewer as well as their “performative gestures” interacting within the entire process.

The interactive book “la_alma” documents a remote history when viewers are invited to touch a hand-made book. By doing so, they trigger live video edits and “converse” with video projections on blank pages. Montage imagery of ticking machines, microfiche, obsolete handwriting, and the palimpsestual nature of the former East Berlin all reveal memories of a woman who has nearly forgotten language. As the narration describes, speaking with Alma during her Alzheimers was like “jumping through a word search for remember.” Using the experience of conversing with a woman who accesses language with no chronology or logical links as a model, “la_alma” accesses randomly generated movies through the touch of a blank book. Movies are derived from Alma's forgotten 80 letters and documents from Berlin to Brooklyn in the 1860s to 1880s. Letters were scanned and addresses were traced and videotaped to create a seemingly random logic in the content of the “la_alma” narrative.

research.umbc.edu/~lmoren/

Special thanks to Gilles Alliume for his production assistance and Tom Demeyer for programming Image/ine at STEIM, Amsterdam.

1. Alexander, C. Bernstein, C. Higgins, D., Drucker, J., Knowles, A. & others. (Eds.). (1995). *Talking the boundless book: Art, language and the book arts*. Minnesota Center for Book Arts.
2. Udagawa, M. & Moeslinger, Z. (founders of Antenna Design, New York), Visiting Lecturers, University of Maryland Baltimore County, November 1999.



“la_alma” installation view (detail).

Installation view.



Still frame from pages of interactive book.

Leveraging the Team Across the World: Creature Development for “Pitch Black”

Paul Franklin
Double Negative Ltd.
paul@dneg.com

Elie Jamaa
Julian Mann
Double Negative Ltd.

Digital creature animation can be costly. The “Pitch Black” script called for over 100 photo-real “critter” animations. The entire production, however, was budgeted at a modest US\$25 million. Filming was done in Australia, which offered both favourable exchange rates and exotic sci-fi landscapes.

Double Negative, based in London, was formed specifically to bring the “Pitch Black” creatures to life. Rather than wait for shooting to finish, the company became involved in production immediately. A small, four-person CGI team was assembled in London and shipped to the production’s Australian studio complex.

Being on-site during filming gave the team great access to the director and VFX supervisor, but time and budget restrictions meant that this Australian team would be responsible for all character setup. Furthermore, the final total team size would be only 20 people at the height of post-production. Most of this team had yet to be recruited, as Double Negative was still constructing its London facility base.

Rather than create tools from scratch, the team chose off-the-shelf applications wherever possible. Alias|Wavefront’s newly released Maya was chosen as the pipeline’s backbone. Maya’s scripting capabilities were brought into play; whenever any task, whether complex or trivial, was performed, it was then scripted so that repetition was greatly simplified. Starting with character setup, this philosophy was applied to the entire CGI process.

The team used time-zone differences between Australia and the UK to their advantage. Tools and setups developed during daytime in

Australia were emailed to London for testing overnight. A whole day’s feedback and comment would then be waiting in the inbox the following morning.

A tight working relationship developed in the Australian group. All team members involved themselves in technical and character-based issues. This combined TD/ animator approach produced a character setup that automated as much as possible whilst maintaining a clear, intuitive interface. New recruits with little or no experience with Maya were thus able to quickly start animation.

The pipeline effectively leveraged the individual skills of team members, encapsulating technical specialities in an easy-to-use tool set. Ultimately, each artist could take a shot through with minimum technical support from the rest of the team.

After filming ended, the Australian group rejoined the London character team. The production team returned to Los Angeles for editing and dubbing. Continuing the concept of working at a distance, Double Negative constructed a bespoke ISDN dailies transfer system (again largely from off-the-shelf components). Shots completed during the day in the UK were uploaded to the US for viewing by the director upon his arrival at the LA office each morning.

By approaching each aspect of the character process as part of a whole unified by the latest tools, the “Pitch Black” team showed that a small group divided by great distance can produce substantial amounts of quality creature animation in a short space of time.



Detail of “Pitch Black” Creature.



Uh-oh.

There are many different technologies for producing line art. Analytical methods, or object-space methods, can give very complete line information but are not as easy to implement as scan-conversion methods.¹ Hardware techniques are very fast, but may render with aliasing, and they may not render crease lines.² Using edge detection on geometry maps has become very popular, but it also requires supersampling and filtering to avoid aliasing.³

The A-buffer is a robust anti-aliasing hidden-surface evaluator.⁴ The technology uses a bitmask as an occlusion map for each sample scanned into each pixel. A typical A-buffer coverage mask covers the area of one pixel. Let us define an extended coverage mask as one that covers an area larger than one pixel. For our purposes, we need a mask that extends by one sub-pixel row or column into the area of each neighboring pixel. The extended coverage mask resembles a convolution region of interest and could be used for filtering. See Figure 1.

With such masks, we can conduct line detection between samples in the same pixel. Without the extended coverage information, iteration over neighboring pixels would be necessary to detect lines occurring at the pixel borders. Figure 2 demonstrates line detection between two 16-bit extended coverage masks. Figure 3 gives C code for executing these functions using bitwise operations.

Modifying an A-buffer to produce this information is not particularly difficult. We benefit from associating line samples with coverage masks because some subsystems in the rendering pipeline can process line samples as if they were surface samples. However, a generalized line-shading architecture varies greatly from a surface-shading

architecture. The look of a line can rarely be decided by point-sampling. Aesthetic choices are often made based on local area information in world space or image space.^{5,6,7} Artists may also want to modify line widths or placement to produce a more clear illustration. Some shading features may even involve repositioning line samples after scan conversion.

Given the development effort, the A-buffer can cope with many of the demands of a line-shading architecture. It can render anti-aliased surfaces and lines in one pass with transparency. Polygons need not be convex or preprocessed. Scan conversion information can easily be provided to subsystems internally for accurate line detection and shading. These features make the technique applicable to a wide range of line-art rendering projects.

References

1. Elber, G. & Cohen, E. (1990). Hidden curve removal for free form surfaces. *ACM SIGGRAPH 90 Conference Proceedings*.
2. Raskar, R. & Cohen, M. (1999, April). Image precision silhouette edges. *ACM Symposium on Interactive 3D Graphics*.
3. Saito, T. & Tokiichiro, T. (1990). Comprehensible rendering of 3D shapes. *ACM SIGGRAPH 90 Conference Proceedings*.
4. Carpenter, L. (1984). The A-buffer, an antialiased hidden surface method. *ACM SIGGRAPH 84 Conference Proceedings*.
5. Dooley, D. & Cohen, M. (1990, March). Automatic illustration of 3D geometric models: Lines. *ACM Symposium on Interactive 3D Graphics*.
6. Kamada, T. & Kawai, S. (1987). An enhanced treatment of hidden lines. *ACM Transactions on Graphics*, 6 (4).
7. Gooch et. al. (1999, April). Interactive technical illustration. *ACM Symposium on Interactive 3D Graphics*.

Liquid Time is an artistic investigation into new forms of representation made possible by digital technology. The goal of the project is to display multiple times or perspectives simultaneously.

In the Liquid Time installation, the motion of a viewer's body in the room disturbs time and space in a video sequence, causing the sequence to fracture and ripple. As participants move closer to the projection screen, they push deeper into time, but only in the area of the screen directly in front of them. As they move away, the image heals behind them like a pond returning to stillness.

The fragmentation of video imagery in the Liquid Time installation is accomplished through video processing software written by the artist. The software works by fragmenting data from a stored digital video file into units smaller than the standard video frame. Vertical slices from different frames of the same video are presented and played simultaneously, creating the effect of video cubism.

This “deconstruction” of the video frame as a unit of display is only possible when a video sequence is stored and manipulated digitally. In a traditional linear video sequence, relationships are formed only between images in consecutive frames of the video. In the Liquid Time installation, relationships are formed between countless image fragments from many areas of many frames. The process creates a surprising visual richness — even within a seemingly mundane video sequence.

While the description of the fragmentation process in Liquid Time sounds chaotic, the process actually creates imagery similar to that of a reflection seen in ripples of water. Users advance time in a certain area of the screen by moving closer to the screen, but when they move from the area, the disturbed areas of the screen return smoothly to the beginning of the video sequence, passing through all the frames (and moments of time) in between. Viewers who walk through the installation from one side of the screen to the other essentially create a “wake” in time that heals smoothly behind them as they pass.

Participants' locations in the Liquid Time installation space are tracked by an overhead video camera. Using live-camera feed as the input mechanism not only allows users a hands-free mode of interaction, but also allows for many areas in the stored video file to be manipulated simultaneously.

Tracked points in the incoming camera feed correspond to multiple users, or multiple points on one user's body, as opposed to the single location point of a mouse. These tracked points are used to manipulate the video fragments on an area of the screen in front of the participant, and to adjust the fragment to the moment in the time sequence corresponding to the participant's proximity to the screen. Since user's bodies viewed from the overhead camera are usually smooth in their contours, the fragmentation of time in the stored video reflects the gentle curves of human bodies.

The Liquid Time installation currently cycles through a series of video clips. The imagery consists of cityscapes and urban street scenes. A variety of filming techniques and camera moves are used to explore the different visual effects created when a sequence is fragmented. Sequences with camera pans and zooms allow for space, as well as time, to be fragmented and manipulated by participants. While the video clips used in the current version of the installation have been thematically selected, other video sequences can easily be loaded into the program and explored.

The Liquid Time video processing software has become a tool for personal artistic exploration, as well as the basis for the current Liquid Time installation.



Installation view.



Screen shot from installation.

The cost of building a digitizing system for light-field³ or Lumigraph² rendering can be prohibitive. Furthermore, the physical size, weight, and complexity of these systems has, in effect, limited their use to small objects and indoor scenes. The primary motivations for this project are to reduce the acquisition cost of light-field capture devices and create a portable system suitable for acquiring outdoor scenes. We have designed such an apparatus using readily available parts. One of our strategies for reducing system cost has been to rely on software to correct as many of the geometric and photometric inaccuracies as possible. Our resulting light-field acquisition device can be built for under \$200.

Our light-field acquisition system employs a modified low-cost flatbed scanner, an array of plastic lenses, and a 12-volt battery (see Figure 1). The scanner is interfaced to a standard desktop PC for indoor use or a laptop for outdoor experiments. Our prototype uses a parallel port scanner, the UMAX Astra 2000P (street price \$70), with a 600x1200 dpi optical resolution at 36 bits of color depth. We disabled the internal fluorescent bulb so that the all of the captured luminance would be attributable to the attached lens system. Focused onto the glass of the scanner is an 8-by-11-grid assembly of one-inch plastic lenses that we found on toy “bug boxes” (plastic boxes used to store and display mounted insects, \$0.50 each). To make the scanner mobile, we replaced its DC power supply with a 12-volt lead acid battery.

Although construction of our light-field capture system is simple, the most significant challenges involved processing the raw scanner output (see Figure 2 for a 2x2 subset of the raw scan). The low quality of these raw images is due to several factors. Since we have used short-focal-length single-lens optics, each individual image exhibits considerable radial distortion. Furthermore, the unknown intrinsic and extrinsic

camera parameters for each lens in the array must also be computed. However, this calibration can be computed once and then reused for multiple scans. In addition, it is necessary to correct the color balance of the raw images to match external lighting conditions. We assume that the built-in color balance of the scanner has been tailored to the spectrum of the original fluorescent light that we disabled. For the most part, we used off-the-shelf image-processing software to correct these errors (see Figure 2).

RESULTS

We have captured various scans of both indoor and outdoor scenes. The image resolution of each image in the 8-by-11 light field is about 600 x 600 pixels. However, the usable region of each image is about 590 pixels in diameter, because of the circular lenses.

CONCLUSION

The primary emphasis for this project was low cost. We have developed an inexpensive, portable system with a price tag under \$200. We hope that availability of an inexpensive light-field capture device will provide more image-based rendering researchers with access to real-world data sets.

References

1. Devernay & Faugeras. (1995). Automatic calibration and removal of distortion from scenes of structured environments. In *SPIE* 2567, 62-72.
2. Gortler, S.J., Grzeszczuk, R., Szeliski, R. & Cohen, M.F. (1996). The lumigraph. *Proc. of ACM SIGGRAPH 96*, 43-54.
3. Levoy, M. & Hanrahan, P. Light field rendering. *Proc. of ACM SIGGRAPH 96*, 31-42.
4. Tsai, R. Y. (1986). An efficient and accurate camera calibration technique for 3-D machine vision. In *CVPR*, 364-374.
5. Zhang, Z. Y. (1999). Flexible camera calibration by viewing a plane from unknown orientations. In *ICCV*, 666-673.



Figure 1. The capture device.



Figure 2. (top) 2x2 subset of the raw scan. (bottom) 2x2 subset after color correction.

MAJOR DAMAGE: The Virtual Garage Studio Processes for Producing Animation

Demolition Crew
Chris Bailey
Creator/Director

Kellie-Bea Cooper
Producer
kellie-bea@kellie-bea.com

Doug Cooper
Computer Graphics
Supervisor

The character-animation industry is filled with talented and brilliant people with one goal in common: to create and share good storytelling. As hardware and software becomes more affordable, there will always be three common challenges that make up the core of any animation process: the art, the technology, and the management. These challenges can be exponentially multiplied when undertaking an independent production.

Often professionals in our industry are looking for a creative outlet or an experimental test bed outside of their daily grind. The drive and the passion is strong, but the problems facing them are significant. With no monetary support and little time outside their already-busy workload, they must find new and creative ways to accomplish what large, well-funded studios are able to do.

Working after hours, with no financial subsidies, in an industry of limited time and support for the exploration and experimental project, this team has beaten the odds, manufactured the moxie, and boldly gone where not many underground, secret-handshake and decoder ring club houses have gone before.

The demolition crew on MAJOR DAMAGE shows their techniques and strategies to prove the success of the Virtual Garage studio.

SOLUTION: KEEPING COSTS DOWN

- We used a production Web site to exchange models and scene files with all the members of the crew spread throughout the city. This was possible through the use of commodity PC hardware and freely available software like Linux, MySQL, Apache, and SaMBa to build a robust, low cost, Internet-enabled server platform.
- Though Maya referencing is commonly used to share files within a centralized studio environment, we were able to leverage this technique to create incredibly small scene files that could be transferred over the Internet quickly.
- MEL scripting techniques were devised to allow multiple artists to collaborate on construction and rigging of a single character simultaneously while working from different locations.
- Careful design of our character controls allowed us to begin animating even before the characters had been completely built. Skeletons were built to the final model, and low-resolution proxy models were constructed for animation. Animation then began immediately, while final skinning and musculature were still being applied to the high-resolution model.

www.major-damage.com



MAJOR DAMAGE, Chris Bailey



Merging CG objects onto video sequences is an important technique for many applications such as special effects in movies and augmented reality. Automatic registration of virtual objects with video images is also called “match moving.” However, conventional match moving techniques mainly focus on rigid movement of single objects. Merging virtual objects with articulated figures such as human bodies is still a difficult problem. In this sketch, we propose a new technique for automatic registration of virtual objects with human body images.

First, we estimate a 3D pose of a human body from video. We represent the human body by an articulated structure consisting of 10 rigid parts corresponding to body, head, upper arms (ur-arm, ul-arm), under arms (ul-arm, ur-arm), upper legs (ur-leg, ul-leg), and under legs (lr-leg, ll-leg). The model has a tree structure with a root at the trunk, and it has a local coordinate system whose origin is located at a joint (Figure 1). The motion parameters are estimated using the spatial and temporal gradient method.¹

Next, we generate CG objects using the pose parameters of the human body. The self occlusion due to the body parts is also estimated using the pose parameters and the 3D shape of the body. The CG image can be composed with the input image using the alpha channel.

In Figure 2, we show the result of merging CG cloth onto a human video sequence: (a) is a typical frame of the input image sequences; (b) is the tracking result; the wireframe of the human body model is overlaid onto the input image; (c) is the generated cloth image using 3D body pose; and (d) is the merged result.

A possible application of the proposed technique may be the virtual fashion simulator in Figure 3. A person can wear any virtual cloth before a virtual mirror. More experimental results and technical details can be seen at: www.vision.info.eng.niigata-u.ac.jp/en/

Acknowledgment: The authors would like to thank Yumi Hoshino for her human illustrations.

Reference

1. Yamamoto, M., Sato, A., & Kawada, S. (198). Incremental tracking of human action from multiple views. *Proc. of IEEE CVPR98*, 2-7.




Figure 1: Relationship of camera coordinate system, world coordinate system, and body model.



Figure 2: Example of merging CG cloth onto a human video sequence.



Figure 3: Virtual fashion simulator.

Computing the Euclidean distance from a point to a complex polygonal shape is a fundamental problem in computer graphics. Distance carries more information than occurrence or non-occurrence of collision because it permits prediction, use of coherence, and dynamic path modification.




Figure 1




Figure 2

Our algorithm uses a multi-resolution hierarchy of bounding volumes generated by geometric simplification. This hierarchy can be used for efficient display as well. As the mesh refines, improving closest point estimates appear to sweep the mesh until the true closest point is found; hence the name of the algorithm (Figure 2). The refinement process can be interrupted at any time, providing an approximation to the distance. Otherwise, the final result is the exact closest point and distance, within floating-point precision.

Our algorithm is dynamic, exploiting full coherence between subsequent queries (without caching the closest point). Our dynamic method also applies to a mesh that deforms, as illustrated with a simulation of particle dynamics (see Figure 1). Achieving from about 500 to several thousand queries per second on complex polygonal shapes (50,000 triangles), our method has applications both for interactive and photo-realistic graphics.

BOUNDING ELEMENTS

While various types of bounding elements could be accommodated (axis- or shape-aligned rectilinear boxes are a common choice), we have chosen as a bounding element the set of points in space whose distance to a shape portion is less than ϵ . In the plane, the elements are formed with segment-aligned rectangles capped with two half-circles.

Referring to Figure 3, the distance d_0 from p to the portion of the shape S contained in a bounding region R is such that $d - \epsilon \leq d_0 \leq d + \epsilon$, where d is the distance to the shape portion defining R .

PRIORITY PROCESS

After computing the distances to all regions in the coarsest level of the hierarchy, each region is associated with an interval of the type $[d - \epsilon, d + \epsilon]$ and indexed in a priority queue. When the key used in the queue to assign priorities is $d - \epsilon$, the interval listed in the front of the queue represents a minimum bound to the Euclidean distance between the query point and the shape. Figure 3 shows this option, where the front of the queue is on the left side, and the back on the right side. The $d + \epsilon$ value corresponding to the leftmost interval provides an upper bound to the distance. This upper bound can be used to prune all the elements of the queue whose key exceeds the bound.




Figure 3

SMALL MOTION OF THE MESH OR QUERY POINT

Each interval in the queue is expanded on both sides by an amount equal to the magnitude δ of the motion. In this way, we have a correct bound on the distance, however conservative. The priorities are not affected, because each interval is expanded by the same amount. The intervals are expanded all at once, by-passing the priority queue mechanism and the logarithmic costs involved.

NON-RIGID MOTION OF THE MESH

When comparing previous bounding volume hierarchies with an error-bounding mesh simplification, mesh simplification computes bounding volumes locally, in a bottom-up fashion, instead of a top-down fashion. For a smooth deformation, the motion of the mesh can be well approximated locally by a rigid transformation. Hence our algorithm will produce almost the same bounding volumes, albeit rigidly moved. This technique was used to produce collision detection animations with a mesh that deforms.



Figure 4

A Method Of Generating Pavement Textures Using The Square Packing Technique

Kazunori Miyata
Tokyo Institute of
Polytechnics
miyata@img.t-kougei.ac.jp

Takayuki Itoh
Tokyo Research Lab,
IBM Japan

Kenji Shimada
Carnegie Mellon University

The aim of this method is to generate pavement textures in which the paving stones are mostly of equal size and may form a “stream.” The packing pattern for a given area is generated by calculating the equation between the forces of scattered particles, and the pavement texture is obtained by generating stone textures for each packing area. In this method, the boundary shape of a pavement and the average size of stones are specified as geometry data, and the roughness value, colors, and optical attributes of stones are given as attribute data.

SQUARE PACKING FOR QUADRILATERAL MESH GENERATION

Square packing¹ is a method for packing square particles that have attractive or repulsive forces such as the van der Waals force. The method assumes that each particle has a potential field of forces as shown in Figure 1. Given a 2D domain and a vector field, the method generates a set of the particles that is well-aligned along the vector field, by solving an equation of motion for the particles. Figure 2 shows an example of the result obtained by the method.



Figure 1: Potential field of a square particle. Left: 3D representation. Right: 2D isocontour representation.




Figure 2: Square packing process. Left: input 2D domain and vector field. Center: initial particle locations. Right: final particle locations

STONE TEXTURE GENERATION

The pavement texture is obtained by generating stone textures for each cell that is the individual area generated by the square packing method. Each stone texture is generated in three steps: generation of basic stone meshes, smoothing, and displacement. First, basic stone meshes are generated according to the geometry of each cell. The 3D shape is obtained by sweeping the base, which corresponds to each cell, by the height of a stone. The basic stone mesh is then generated by deforming this cube, displacing each top corner randomly. After basic stone meshes have been generated, they are subdivided by using Loop's subdivision method,² which refines the given meshes. Finally, the stone meshes are displaced by means of a fractal noise function, which gives the surface bumps. The amplitudes of a fractal noise function can be adjusted to create various stone textures.

RESULTS

The processing time depends on the number of the stones. It is about five seconds on average for the square packing process and about 10 seconds on average for stone texture generation (on an Intel Pentium2 300 MHz processor, and for a texture size of 512 by 512). Figure 3 shows examples of generated pavement textures, and Figure 4 shows a CG image that uses a pavement texture.

References

1. Shimada K., Liao J., & Itoh T. (1998). Quadrilateral meshing with directionality control through the packing of square cells. 7th Intl. Meshing Roundtable, 61-76.
2. Charles Loop. (1994). Smooth spline surfaces over irregular meshes, *Proc. of SIGGRAPH 94*, 303-310.





Figure 3: Examples of pavement textures





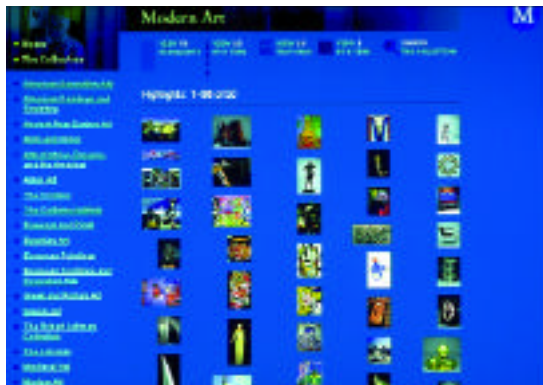
Figure 4: Sample CG Image

Web sites themselves aren't so new anymore. Most organizations, institutions, companies, large and small businesses, and many individuals have them. What makes a Web site unique, exciting, or "of note" is that it "works." It does something that is beneficial to us by making our lives easier, more enjoyable, or more efficient. Many times these delights come from new technology, an interesting conceptual direction, or even a low-tech idea. The development team at Icon Nicholson and the Metropolitan Museum of Art feel that we have achieved these goals in some areas of the recently launched redesign of the Metropolitan Museum of Art's Web site (www.metmuseum.org).

The museum itself is a special place to visit. As you ascend the grand staircase from Fifth Avenue in New York City into the arched great hall of the museum and enter the halls, you are surrounded by people from all over the world admiring great paintings and sculpture that span 5,000 years. Our task was to transform that experience to a computer interface, a much more intimate situation: one person interacting with

a computer screen 14 inches away. We strove to maintain the classic, sophisticated feel of the museum while giving the online user a unique experience. The visitor to the museum is treated to the awesome grandeur of the staircase and the great hall, while the online visitor is treated to a different piece of art from the collection on the first screen every day accompanied by an "art fact" or curator observation. Each of these welcomes is special but unique to the approach, whether it is virtual or real.

Additionally, the online user has access to more information about each piece of art — from the provenance to artists' biographies, to information on signatures, marks, and inscriptions, for instance. The Web site includes more than 3,000 works of art from the Museum's collection. The visitor to the museum has the benefit of seeing the actual work of art as opposed to the digital representation of it. We didn't want the Web site to replace a visit, or a visit to deter a user from going to the Web site. They are different experiences, each with unique benefits.



The input of 3D data (whether positions, vectors, or orientations) is often essential in order to interact with 3D objects. However, on a simple monitor, and using 2D devices like the mouse, it can still be a tedious task. These devices are well suited to 2D interaction, allowing any visible target to be pointed at with a reasonable precision. Moreover, in 2D drawing software, it is not uncommon to see a grid that offers some help in estimating distances or even angles.

We were interested in extending these concepts to the 3D case and have implemented a widget that combines the functionalities of 3D movement (to position points in space, for example) with measuring abilities inspired from the 2D grids. The Metric Cursor was primarily intended to be used in a modeling system. This 3D cursor, appearing as three mutually orthogonal planes, can move freely in the scene and is used to input 3D points.

Essentially manipulated with mouse and keyboard (for the moment), the cursor is moved by successively translating it along one of its axes.² This method has drawbacks, but a rotation of the scene quickly solves problems due to the superposition of two axes and still allows free movement in space. When only one face of the cursor is displayed, its movement is restricted to the 3D plane supporting it. The cursor can be rotated around any of its three axes. Reorienting it simplifies displacement when the privileged directions of translation are known. This allows the cursor to move easily on a face. To speed up manipulation of the cursor, the user is able to align it with an edge (one axis will be made parallel to it) or two edges (one face of the cursor will be orthogonal to the plane defined by the edges, if any) and make the cursor jump to a selected vertex. Thus, this is somehow equivalent to snap-dragging¹ without really implementing it, at least for scenes made up of only polygons.

Two types of grids are linked to the cursor:

- A regular 3D grid, an extension to three dimensions of the grids found in 2D drawing software. Only displayed on each face of the cursor, it can help in estimating distances and is also used to constrain the movements of the cursor by snapping its center to the grid's nodes.

- Three radial grids, one for each axis of the cursor, displayed on the face orthogonal to the axis (see Figure 1). These assist in controlling the angle of rotation in the same fashion as the previous grid (by snapping to the nearest orientation).

Using the translation of the cursor, one can measure the distance between any two points. At the beginning of a translation, a mark is set to recall the initial position, and the distance between the cursor and this mark is displayed as the actual euclidian distance, and as the number of steps taken on each axis of the cursor.



The step size of these grids is, of course, variable, allowing, for example, the user to measure a distance and make it, or a fraction of it, the new step size for the 3D grid. In addition to being displayed on each face of the cursor, the 3D grid is also mapped, during a translation, onto the "displacement trace" (see Figure 1), which consists of three perpendicular faces joining the initial position of the cursor to its current position. This makes evaluating distances easier during a cursor movement.

This widget aims at supporting 3D drawing and measuring, placing vertices in space precisely as required. It does a good job in its ruler and protractor tasks, especially for distances between two points that do not belong to the same edge. The same could be done to retrieve the angle between two non-connected edges. Of course, the 3D data obtained from the manipulation of the cursor can become input to some complex modeling functions.

After a usability evaluation, the shape of the Metric Cursor might evolve to reduce scene clutter, which is a problem that may not be as important in a virtual reality environment such as on the Responsive Workbench. However in this type of environment, manipulation of the cursor should benefit from the use of 3D input devices. Finally, to modify the cursor in order to support the modeling of more general shapes than polygonal ones, it might be necessary to have recourse to other types of grids.

References

1. Bier, E.A. (1990). Snap-dragging in three dimensions. *Computer Graphics (1990 Symposium on Interactive 3D Graphics)*, 24 (2), 193-204.
2. Nielson, G.M. & Olsen, D.R., Jr. (1986). Direct manipulation techniques for 3D objects using 2D locator devices. *Proceedings of 1986 Workshop on Interactive 3D Graphics*, 175-182.



Production of stage events such as plays and concerts is a time- and money-consuming process. Beneath the huge technical effort, design is a creative and collaborative process. Before the final stage setup is completed, directors, stage technicians, etc. are evaluating different arrangements, often changing the complete setup. During this experimentation phase, 4:1 stage models have already proven to significantly reduce the effort in terms of time and money. While the downscaled stage can create realistic light, sound, and fog effects, it lacks support for choreography. Pure VR simulations are not applicable due to the complexity of lighting conditions and their non-intuitive user interfaces. To bridge the gap between intuitive usage of a real model stage and the power of computer visualization, we have developed the Mixed Reality Stage. This approach uses augmented reality¹ to enhance a real stage model with computer-generated animated 3D objects. It aims to provide the basic mechanisms required to evaluate stage setups while reducing the need for rehearsals.

MIXED REALITY STAGE ENVIRONMENT

The Mixed Reality Stage consists of a real model stage and its virtual counterpart. While the real stage simulates light and sound effects, the virtual stage enhances the real model with animated characters and additional requisites. Our augmented reality system allows us to provide users with an integrated view of the real and virtual stages. The current prototype is based on a standard PC with stereo-capable graphics accelerator to provide an affordable solution. The system uses an InterSense 6DOF tracking system for viewpoint tracking and light-weight, see-through glasses for projecting the virtual 3D objects onto the real model stage. Users are free to move around the model stage. Multiple users can view the model stage simultaneously due to the distributed architecture of our underlying VR toolkit.² While this scenario also provides support for modifying materials of real stage objects, it is especially useful for simulating the movement of characters such as actors, dancers, or musicians.

One of the major advantages of the real model stage is the highly realistic simulation of light effects. To guarantee similar lighting conditions for real and virtual objects, stage lights additionally have to be modeled

within the virtual scene. Stages often contain 50-200 individual light sources, which cannot be rendered in real time on PC-based graphic systems. Two mechanisms can be applied to reduce the complexity of lighting conditions. First, many stage lights are not applied individually, but rather as light bars consisting of six or eight spot lights. During rendering, we approximate each bar with a single light source. Second, each object on the stage usually receives most of its light from particular light sources. Thus, appropriate scene graph arrangements can be used to ensure the availability of the required light sources.

Shadows of virtual objects cannot be displayed with the current see-through technology. Since dark objects always appear transparent to the user, it is not possible to superimpose real objects with dark virtual objects. We model all real objects of the stage as black, unlit virtual objects. This guarantees that the occlusion of virtual objects by real objects is correctly visualized. Otherwise, distant virtual objects could occlude close real objects, destroying the illusion of augmentation.³ Obviously, the locations of these virtual objects have to be synchronized with their real counterparts.

FUTURE WORK

In our future work, we will provide seamless integration and synchronization of the real lights into the AR system. We will realize additional tools based on tangible interfaces to support interactive compositions of choreography. This approach will be based on downscaled actor dolls, which will be used as tangible interfaces. A sourceless and wireless user-tracking system as well as an advanced camera-based tracking system will be developed.

orgwis.gmd.de/projects/CAMELOT

References

1. Azuma, R. (1995). A survey of augmented reality. *Presence*, 6 (4), Barfield, W., Feiner, S., Furness III, T., & Hirose, M. (Eds.), 355-385. Cambridge: MIT Press.
2. Broll, W. (1998). SmallTool - A toolkit for realizing shared virtual environments on the internet. *Distributed Systems Engineering, Special Issue on Distributed Virtual Environments*, 5, 118-128.
3. Wloka, M. & Anderson, B. (1995). Resolving occlusion in augmented reality. *Proc. of the ACM Symposium on Interactive 3D Graphics*, 5-12.



The Mixed Reality Stage: mixing real and virtual objects on a downscaled model stage.

For "The Hollowman," Paul Verhoeven's twisted revision of the "Invisible Man" story, Sony Pictures Imageworks took on the challenging task of bringing the director's vision to the screen in what will be seen as a monumental achievement in visual effects.

The film requires the actor (Kevin Bacon) to transform from visible to invisible in various stages. Sony had to come up with a way to visualize the musculature, vascular system, skin, and physical performance of an actor as well known as Kevin Bacon. Most shows featuring digital actors only have to deal with the complexities of matching that actor's exterior look. The task of constructing the skin became the first of many computer models needed for the Hollowman digital character.

One of our challenges was the level of anatomical and physiological study necessary for the task. It was obvious from the beginning that since humans and apes have such similar biology that it would make sense to share some of the geometry between the two. Of course, each element would have to be resized to fit the other's proportions, but it was faster than building everything from scratch twice.

The first step was to build the digital characters from the outside in, so Kevin Bacon was cyber-scanned in various poses. This gave us polygonal geometry from which we could make a NURBS skin. Scanning was far more accurate than sculpting a clay model of Kevin. Since this is the first thing people will see of our digital Kevin in the movie, it was very important that the sculpture and the real Kevin be identical. Scanned data can present undesirable results in some areas, so a great deal of work went into fine detailing for a perfect match.

Once the skin was approved, it was immediately obvious that we had another challenge on our hands. When modeling dinosaurs, space creatures, or animals for most movies, you can get away with making them symmetrical. However, real people are not, so for the digital Kevin, we couldn't just build one side and make a flip copy for the other. The entire model had to be sculpted, and that would, of course, require more time.

Before we started on the insides of the character, physiology experts assisted with the breakdown of both Kevin's and the gorilla's anatomies. This helped tremendously because the tight production schedule didn't allow for us to get something like an MRI of Kevin. We had to get as much information as we could from Kevin's skin. He had to be measured and photographed in various stages of physical stress if we were to mimic his physique correctly.

We then laid the foundations by sculpting a complete skeleton constrained by Kevin's proportions. Here, the modelers had to work closely with what is known as the physiquing department, the group that creates the controls that enable modeled geometry to move with aesthetic accuracy. That meant we had to sculpt not only for physical appearance, but also for functionality.

The calcium bone placement was extremely important, because the muscles have to fit and anchor correctly off the bones, and because joint placement determines how well the digital actor would move and mimic the real Kevin Bacon. Placement of the spinal column and the

skull was a challenge in itself. A cast was made of Kevin's teeth, and our digital replica was placed correctly using various references. However, only tell-tale signs in Kevin's skin could be used for the rest of the skull. Some scenes needed a skinless, muscleman version of a walking, talking Kevin. For Verhoeven, it was very important that the audience recognizes the skull and facial muscles to be those of Kevin Bacon, so intense design work was devoted to the face.

Once we locked down Kevin's skeletal structure, we moved on to the musculature. For the transformation method we were using, each and every muscle had to be an enclosed volume. This meant that even very thin muscles had to have some level of thickness to them for the volume rendering to work correctly. Sculpting as many muscles as possible within the human body soon filled every air pocket within the skin model. Kevin's muscular definition helped to line up the digital muscles created underneath. Then a plausible amount of space was left for procedural fat and other bodily tissue.

One problem with sculpting tightly packed muscles is the possibility of interpenetration. At some stage or another, all muscles will be seen, so a clash of surfaces had to be avoided. The modelers did as well as they could, but the sheer volume of modeling made it almost impossible to avoid. The software department created plug-in programs for Maya that eased the problem. Alias|Wavefront's Maya was the primary sculpting package, and the plug-in written for it was designed to pull two interpenetrating surfaces apart by using one as a deformer. This was a big plus.

While the muscles were being built, the modelers were careful to leave adequate space for the internal organs. Again, placement of the skeleton was crucial for organs such as the lungs, diaphragm, and heart. An unprecedented amount of detail went into the heart model, since it is such an active organ. Its surfaces had to hold together as it was animated through its beating pattern. A full set of heart-surface veins and system circulation had to be modeled as well.

Kevin Bacon's digital cousin had some close-up scenes in which his skin dissolves to reveal all that is underneath. Sony programmers worked with the modelers to create a new facial animation system to handle the realism needed. Many face shapes were modeled to match various Kevin expressions and then manipulated in a very efficient manner using the new facial editor. Unlike technology we had used previously, the editor allowed control of not only the face shapes, but of Kevin's iris dilation and jaw movement. Phonemes and their jaw rotations could be blocked very quickly, and the whole face animation could be recorded and played in real-time 3D.

Another benefit of the face setup is that the skin shapes automatically controlled the position of the musculature beneath. In real life, it is the opposite, but here the animators only had to look at the skin performance when manipulating expressions for the muscle face.

In the end, a collaboration of talented people managed to create a walking, talking digital actor with a heart, and lungs, and brains, and bones...

With the boom in the e-commerce industry, it is expected that product advertisements will become more interactive in nature. There are a few key techniques in providing such solutions. The first and the foremost solution is VRML representation of the advertised object's 3D model. The other solution is similar to the Quicktime VR technique, where a large number of pictures of the object are stitched together to create a movie. Interactive browsers are designed to play back the movie.

In e-commerce applications, e-shop personnel may need to "scan" thousands of items in the shop within a limited period of time. Data collection is the biggest bottleneck associated with the two approaches mentioned above, since they require either precise movements of camera around the object, or costly devices to create the models/movies. Also, from the bandwidth point of view, such methods are expensive, especially when the e-customer sees the object from only a fraction of the entire view sphere. Hence, most of the downloaded data at the client end remain unseen.

Our novel technique of interactive product advertisement seeks to address the following issues:

- Data collection should be as simple as a few freewheeling "clicks" of the camera around the object.
- The method should not require high bandwidth.
- There should be room for natural objects to be advertised. Our method is based on the image-morphing technique presented by Seitz and Dyer.¹

Our final goal is to give e-customers an interface to look at the object by moving it in two or three possible trajectories around the object. In the data-collection stage, for each of these trajectories, we capture the images of the object from a few key positions. Note that there is no need to follow any strict geometrical relations among these positions, which makes the data-collection process easy. Next, for every "neighboring" image pair in a trajectory, the rectification matrices R_1 and R_2 are computed using the theory discussed Seitz and Dyer.¹ With a user interface, a few matching lines between the two neighboring images are also selected in the processing stage to aid the morphing process. When an e-customer views an object, only two of the images need to be present at the client end. Also, only two rectification matrices corresponding to the image pair and the information of the matched line pairs are required at the client end for computation of the warp and the morph. The client-server communication involves periodic transmission of two still images (in the JPEG format) and a few floating-point numbers corresponding to the rectification matrices and the matched lines. Thus, bandwidth usage is kept minimal. The synthesized image is created by following the three step algorithm in Seitz and Dyer.¹

References

1. Seitz, S. & Dyer, C. R. (1996). View morphing. *Proc. SIGGRAPH 96*, 21-30.
2. Brier, T. & Neely, S. (1992). Feature based image metamorphosis. *Proc. SIGGRAPH 92*, 35-42.

MorphZhapes is an application that addresses the creative opportunities available to any designer, architect, artist, or “maker” who is exploring and learning about the threshold between two- or three-dimensional design, architecture, pure Java programming, and human zeitgeist.

An almost unique characteristic of a work of design or architecture is that it is concurrently dynamic and static. It is dynamic when viewed as a product of the design process, rooted in historical precedents of culture and the arts. The work of art or design is further elasticized by the very nature of its creative components, and by its being the result of manipulating visual or structural entities, which are typically of an elastic character.

A work of design later becomes static when it is frozen at a certain state during the concept’s creative generation, in order for it to be built. In other words, a work of design or architecture is static when viewed through its formalized physical existence. It becomes dynamic when the work is viewed as an instance of a continuum, deriving from the past and projecting into the future.

MorphZhapes seeks to provide the maker, the designer, the architect, or the artist with a flexible, Java-run, cross-platform application that would allow and encourage such experimentation. Being very small in application code size, and cross-platform, MorphZhapes provides an ideal creative application. The maker can morph between two different


shapes (parent A and parent B), each one with a different number of points, and create a third shape, a hybrid child. The miracle of creation is now in a portable, cross-platform Java form.

The aesthetic of this endeavor does not rely much on the final form, but rather in the process, in the intermediate phases these initial shapes go through, as well as in the extrapolations that extend beyond the final form. The user has the capability, through the system, to modify and control the flow of the compositional evolution and replay it many times by varying some or all of the transformational parameters.


MorphZhapes creates a new form of geometry, a liquid geometry, and its results become morphResults. Two buildings being morphed define a new area for architecture, one of morphArchitecture. Two objects morphed would define a new area in design, that of morphDesign. The genealogical trees inherent in building individual traits in human beings now form the basis of the artists’and designers’exploration, leading to a plethora of available mutations. Each object made in this way benefits continually from the character of its parents, while displaying a unique identity.

On a social level, morphZhapes visualizes the myriad possibilities that exist within the genealogical combination between two parents, and it allows its maker to save each liquid 3D iteration between parents and child.

234



A morphed-zhaped process of blending two structures in order to study and experiment further with the in-between states.



A pedigree showing three parents (p1, p2, and p3) and three generations of crossing the morphologies of the two parents and the first generation children.

Motion adaptation is an approach for applying motion-capture data to several characters with different segment lengths and shapes. Although several motion adaptation techniques^{1,2} have been developed to handle segment-length differences between a motion actor and a virtual character, no research has considered the self-intersection problem that may result from shape differences. See Figures 1(a) and 1(b). This sketch considers the self-intersection problem in motion adaptation and presents an approach for avoiding self-intersection in the adapted motion of a virtual character.

OUR APPROACH

Our approach is based on collision detection and avoidance. We first apply the given motion data to the virtual character and check the collisions between segments at each frame. When a collision between two segments is detected, we can avoid the collision in the adapted motion by adding a geometric constraint in the motion-adaptation process that moves the two segments to a collision-free configuration. However, if such geometric constraints are extracted at all frames that contain collisions, the resulting motion may contain jaggedness because collisions at adjacent frames may be avoided in much different configurations. Hence, we gather adjacent frames with collisions into a group and compute a geometric constraint only at the frame where the largest collision occurs.

The geometric constraints are then combined with other constraints and become the input of an existing motion-adaptation algorithm. Since most motion-adaptation algorithms consider inter-frame relationships, a geometric constraint on a frame has an effect on the adjacent frames, and the final adapted motion is expected to have no self-intersection at any frame.

COLLISION DETECTION AND AVOIDANCE

A virtual character usually consists of many segments with complex shapes. Collision detection between such segments is a time-consuming process. To speed up the computation, we simplify the shape of a virtual character with two kinds of simple primitives, spheres and cylinders. See Figure 1(c).

When a collision between two segments is detected, we move one or both of the two segments to achieve a collision-free configuration. To determine which segment should move, we assign a priority to each segment, depending on its mobility. When a segment is moved, the moving direction is selected to minimize the amount of the movement to preserve the original motion characteristics.

EXAMPLE

Figure 2(a) is an adapted motion of a virtual character with a self-intersection. By using the technique proposed in this paper, the self-intersection has been avoided in Figure 2(b).

www.postech.ac.kr/~{misterq,leesy}

References

1. Gleicher, M. (1998). Retargetting motion to new characters. *Computer Graphics (Proc. SIGGRAPH 98)*, 33-42.
2. Lee, J. & Shin, S.Y. (1999). A hierarchical approach to interactive motion editing for human-like figures. *Computer Graphics (Proc. SIGGRAPH 99)*, 39-48.



Figure 1: Motion actor and virtual character.
(a) motion actor, (b) virtual character, (c) simplified character.



Figure 2: Self-intersection avoidance example.
(a) self-intersection, (b) avoided.

A science museum is a communication system, a data space where people go to be immersed in information, in the form of real things in a 3D space. While technology should always be secondary to the artifacts and ideas on display, it can be used in creative and subtle ways to enhance instruction and experience. The software design group of the American Museum of Natural History's exhibition department has been exploring the nexus of information, immersion, and interactivity. These are not traditional computer kiosks.

Some scientific information cannot be conveyed with artifacts. Such is the case with our exhibition: "Epidemic! The World of Infectious Disease." Because it deals with complex concepts and very small things, we use technology to enhance the information and experience. A "microbe forest" combines real and virtual 3D models of microbes, all scientifically accurate. A human shadow is the interface to explore normal accumulations of microbes in the human body. Visitors follow an influenza microbe entering the body and attempt to infect or defend the body.


"Endurance: Shackleton's Legendary Antarctic Expedition" shows a completely different use of technology. Like "Epidemic," it is not an exhibition built around objects. It is the incredible story of an expedition early in the 20th century, and the science behind it. Sir Ernest Shackleton and his crew became trapped on the moving ice floes of Antarctica for 18 months and were finally rescued after Shackleton and five others set out across 800 miles of icy ocean.

We wanted to immerse visitors in the experience of the lifeboat journey and show them what an amazing feat it was to travel in a straight line over 800 miles of ocean in 60-foot-high seas using only rudimentary navigation instruments. We created a realistic, high-resolution simulation of rolling seas and projected it on 12-foot screens around the boat to create an immersive effect. Then we outfitted some World War II-era sextants with sensors and microcontrollers, so visitors can take a real sighting, while an adjacent LCD screen compares the results with Shackleton's actual path.

For the "Fighting Dinosaurs" exhibition, we created a similar immersive experience. We simulated the environment of the Gobi Desert 200 million years ago and re-created a fight between a Velociraptor and a Protoceratops, and their subsequent fossilization in a collapsing sand dune. This is presented in the gallery at approximately actual size. We have also created a virtual fossil dig, using sensor technology to allow visitors (several at once) to uncover virtual fossils on a large surface.

Collectively, these projects show how experience helps turn information into knowledge.

236



Part of the exhibition "Endurance: Shackleton's Legendary Antarctic Expedition" at the American Museum of Natural History. This section features an immersive interactive experience about navigation on the open sea.



Screenshot from navigation interactive shows the startup state of the LCD screen. A real sextant is the only input device, and when a visitor manipulates it, the screen responds accordingly.

The combination of powerful, low-cost computers and video cameras has enabled a new class of human/computer interfaces that are more natural for many applications than traditional computer interfaces (such as keyboard or mouse). By incorporating natural interfaces, applications can provide a more intuitive interactive experience and can become accessible to a wider audience.

A great deal of previous work has explored the possibility of processing video for human/computer interfaces. Numerous references and technical conferences exist in the areas of face and gesture recognition, scene interpretation, pattern recognition, and optical tracking. The work presented in this sketch is focused on natural interfaces for various forms of computer entertainment. We overcome many of the issues associated with real-time video processing through the design of highly targeted interfaces. Robust behavior is achieved by limiting the interaction domain to a narrow, well-defined set of situations. These situations can be divided into two categories: those that make use of prop manipulation and those that do not.

In many cases, the use of props can actually make an interface more natural. In our work, we have created several applications for which the manipulation of props seemed the most appropriate interface. In *Medieval Chamber*, a user can orient a rendered sword, flail, and torch by orienting the corresponding stick-like prop. In *Marionette*, a user can orient a cross-like prop to animate a rendered marionette. In *Planet Explorer*, a user can manipulate a globe-like ball to interact with a rendition of the Earth.

For many applications, no props are required to provide a natural interface. In our work, we have intentionally avoided a general gesture recognition approach; instead, we extract very application-specific measurements from the video. In *Air Guitar*, the repetitive motion of a user playing an air guitar sets audio tempo and triggers sound effects. In *Flying*, a user's gross arm motions control the camera altitude, attitude, and speed of a terrain-traversal rendering. In *Hand Puppets*, a user can use hands to control a pair of rendered hand puppets. All of the applications described above are implemented using only inexpensive, consumer-level hardware (namely, a PlayStation2 and USB webcam).



Hand Puppet



Planet Explorer

The NAVE: Design and Implementation of a Non-Expensive Immersive Virtual Environment

Jarrell Pair
Carlos Jensen
Jeff Wilson
Larry Hodges
Dave Gotz

Georgia Institute of
Technology
jarrell@acm.org

Julian Flores
Universidad de Santiago
de Compostela

As demonstrated by the popularity of the CAVE and products such as the Virtual Workbench, there is a great amount of interest in projected stereoscopic environments as alternatives to head-mounted displays. One of the primary obstacles to their widespread adoption has been their high cost.

The goal of the NAVE project was to design a low-cost, PC-driven, multi-screen, multi-user, stereoscopic, virtual environment with many of the desirable elements of the CAVE at a fraction of its cost. The NAVE was built at a total price of less than \$60,000.

The NAVE is a three-screen environment. Each screen measures eight feet (2.4m) wide and six feet (1.8m) high. The screens are positioned at 120-degree angles to each other, producing a three-sided display area sixteen feet (4.8m) wide and approximately seven feet (2.1m) deep. The user is in a seat positioned at the center of the semi-circle formed by the three screens. Imagery for each screen is generated on a 500MHz Pentium III PC and back-projected in stereo by means of polarized light. To experience the stereoscopic effects, the user wears inexpensive, lightweight polarized glasses. A fourth PC coordinates the three screen-rendering machines and provides directional sound for the NAVE.

In the tradition of the CAVE, the name NAVE is a recursive acronym (NAVE Automatic Virtual Environment). The name is also an acronym for its design goals: a non-expensive automatic virtual environment. Finally, the name has an architectural context. A nave is the central part of a cruciform church building, appropriate for our first application,

the Santiago 2000 Virtual City. The Spanish city, Santiago de Compostela, has since the early Middle Ages been one of the most important pilgrimage sites in the Christian world. It is said to house the remains of the Apostle Saint James, brought from Palestine after his death in 42 AD. Santiago de Compostela was chosen as one of nine cultural capitals for the European Union in the year 2000. To celebrate this event, along with the general euphoria of the new millennium, the Santiago 2000 Project was commissioned to build a computer-generated virtual environment to recreate the area surrounding the city's historic cathedral.


The NAVE and the Santiago 2000 project was a collaborative effort between the Virtual Environments group of the Georgia Institute of Technology and the Technological Research Institute of the University of Santiago de Compostela, Spain.

PROJECTORS, MIRRORS, AND SCREENS

We utilized VREX 2210 stereoscopic projectors, which use a linear polarization scheme and provide a mono resolution of 1024 x 768. Using mirrors, we employed a folded-optics system to achieve the necessary projector-throw distances.

CONCLUSION

The NAVE succeeds in making many of the advantages of the CAVE available at a significantly lower price. Chief among these is the high degree of immersion of a large-screen environment with a large field of view. In addition, the NAVE's extensive audio features serve to heighten this sense of immersion.



NAVE screen configuration.




Screenshot of the Santiago 2000 virtual city, the first NAVE application.

With the advent of virtual reality, surgical simulators have been proposed to supplement medical education. Realization of this goal requires real-time cutting to simulate surgical blades. Previous algorithms have used moving templates,¹ particle systems,² radial projections,³ tetrahedral subdivision,⁴ and Boolean operations.⁵ Unfortunately, these algorithms were restricted to 2D, lost surface area, were constrained geometrically, did not perform in real-time, or did not include a physically based system.

Stressing interactivity, triangle sets with a mass-spring lattice were chosen to simulate elastic deformations. One variation is the use of vector springs, which were developed by Alan Millman at the University of Illinois at Chicago's Electronic Visualization Laboratory. Spring stiffness coefficients and mass distribution were computed based upon adjacent surface area.⁷

For simplicity, the blade's path is assumed to be performed at a slow rate without any high frequencies, and it is assumed that the cutting path can be linearly restructured from the spring intersections. When two valid samples have been determined, a parallelogram is formed through their positions along a common orientation. Spring intersection is determined by a line-plane intersection routine followed by a quadrilateral inverse mapping. When intersected, two new subsprings are created with duplicate intersection points. For a vertex intersection, the spring is copied, and other connecting springs are partitioned between them. Afterwards, each affected triangle is subdivided between intersection pairs. If a quadrilateral is formed, it is subdivided again. Any new interior springs compute their rest vectors through the addition of their neighbors. Then, all adjacent springs and vertices near the cutting boundary update their stiffness coefficients and mass values according to the new triangle configuration.



Cutting sequence: Left, rendered; Right, springs only.

Figure 1 shows a set of consecutive images of the cutting process, while Figure 2 shows an image from our vitrectomy simulator.⁶


When a triangle is partially intersected, a matching interior intersection at the boundary of the parallelogram is computed. The original triangle is then subdivided into three subtriangles. The new spring rest vectors can be determined by extending one subspring through the entire triangle and solving for its sublength. Finally, the subtriangle containing the original intersection is subdivided again to open or close the cutting boundary. If a spring is grazed, one intersection point is generated. Grazed triangles are subdivided in half.

Near real-time physically based cutting was accomplished on 3D polygonal models. Cutting does lag behind the user's current position since the system maximizes the current cutting plane. Users can adjust a parameter for finer resolution. Some polygonal artifacts do occur, since the system rounds to the nearest vertex to avoid generating small springs.

www.bvis.uic.edu/paul/

References

1. Mazura, A. & Seifert, S. (1997). Virtual cutting in medical data. *MMVR Proc.*, 420-429.
2. Szeliski, R. & Tonnesen, D. (1992). Surface modeling with oriented particle systems. *Proc. SIGGRAPH 92*.
3. Pieper, S. et al. (1992). Interactive graphics for plastic surgery. *ACM Interactive 3D Graphics Proc.*, 127-134.
4. Song, G. & Reddy, N. (1995). Tissue cutting in VE. *MMVR Proc.*, 359-364.
5. Delp, S. et al. (1997). Surgical simulation. *Presence*, 6 (2), 147-159.
6. Neumann, P. et al. (1998). Virtual reality vitrectomy simulator. *MICCAI Proc.*, 910-917.
7. Gelder, A. & Wilhelms, J. (1997). Varying spring constant. *SIGGRAPH 97 Visual Proc.*, page 153.



Vitrectomy simulator.

A force-feedback device is an important 3D haptic display, especially applicable to computer-aided design (CAD), where 3D object manipulation is essential. Each haptic rendering¹ requires a high refresh rate (1KHz) and collision detection between the cursor and an object, which can be very time consuming when the object has a complex free-form surface, such as a NURBS surface. For CAD applications, smooth haptic interpolation of polygonized surfaces using a technique called “force shading”^{2,3} is useful for maintaining high haptic-rendering rates by reducing the cost associated with computing the exact collision point.

Obviously, coarser polygonization is better, as long as the interpolated shape is felt as being smooth and not different from the original surface. But how coarse can it be? The polygonization resolution is defined by human sensory thresholds; however, no data are present on the absolute haptic threshold values of surface smoothness presented by a force-feedback device. Therefore, we have conducted experiments designed to study the absolute thresholds of smoothly interpolated curved surface perception.

EXPERIMENTS

Method

Method of limits. Subjects responded “yes” if the stimulus shape (Figure 1) was felt to be a smooth cylindrical surface, or “no” if not. Resolution angle α was varied in one-degree increments.

Equipment

PHANTOM 1.0 (SensAble Technologies)

Parameters

(Experiment A) Reaction force direction was force-shaded and continuous. Force magnitude was calculated by a simple spring model, multiplying surface stiffness and penetration depth to the nearest

tangent plane, and displacement had C^0 continuity. (Experiment B) Non-interpolated condition and linear interpolation of the two force vectors calculated at the two nearest tangent planes were measured.

RESULTS AND DISCUSSION

Exp. A (Figure 2): The absolute threshold angle is linearly defined by curvature ($1/\text{radius}$) and stiffness. The maximum difference d (Figure 1) for the results seems to be independent of radius (Figure 2 b).

Exp. B: Figure 3 compares the results of Exp. B and results of Exp. A of the same radius and stiffness. A statistically significant difference (5% level) exists between the non-interpolated condition and the two interpolated conditions, but not between interpolation methods. Without interpolation, reaction force direction is not continuous on tangent plane borders, which explains the difference. These results imply that when reaction force direction is smoothly shaded, the smoothness threshold does not differ noticeably whether displacement continuity is C^0 or C^∞ ; the smoothness threshold is determined mainly by the maximum difference between the approximated shape and the original.

Although this research is still in its early stages, the presented threshold values should be useful in determining the polygonization resolution for force-shaded smooth haptic rendering of curved surfaces.

References

- Salisbury, K. et al. (1995). Haptic rendering: Programming touch interaction with virtual objects. *ACM Symposium on Interactive 3D Graphics*, 1995.
- Morgenbesser, H. B. & Srinivasan, M.A. (1996). Force shading for haptic shape perception. *ASME Dynamics Systems and Control Division*, 1996.
- Ruspini, D. et al. (1997). The haptic display of complex graphical environments. *ACM SIGGRAPH 97*.




Figure 1: Sectional image of stimulus shape. A cylindrical surface, defined by resolution angle α , is interpolated in different ways. The axis of the cylinder is placed along the Z-axis, or toward the subject. d : Maximum difference between the polygonal approximation and the true surface.




Figure 2: Results of Experiment A.




Figure 3: Comparison of interpolation methods.

The goal of mixed reality (MR) technology is to realize environments that seamlessly integrate both real and virtual worlds. In an MR system, mutual occlusion of real and virtual objects enhances the user's feeling that virtual objects truly exist in the real world. This is an essential feature for some MR applications, such as architectural previewing. In addition, in terms of cognitive psychology, incorrect occlusion confuses users. We propose a novel optical see-through display that can present mutual occlusion correctly.

There are two types of head-mounted displays: video and optical see-through.¹ Video see-through displays can handle the occlusion problem easily if the system knows depth information in a real scene. However, they degrade rich information from the real world in terms of low spatial and temporal resolution, and limited depth of field, fixed focus, and so on. Another demerit is that users may lose their view when the system malfunctions. On the other hand, typical optical see-through displays preserve the real image as it is without any degradation. But they cannot display mutual occlusion correctly, since synthetic objects always appear as semitransparent ghosts floating in front of the real scene. In addition, because it is influenced by the color of the real image, each pixel of a synthetic image never shows its original color.

There have been a few attempts to build new optical see-through displays that attack the occlusion problem.^{2,3} However, all existing approaches are impractical for applications that require a wide working area or outdoor activity, since all of them require special real objects or environmental settings.

OPTICS DESIGN AND IMAGE SYNTHESIS

We propose a novel design of an optical see-through display that overcomes the occlusion problem. The heart of the new design is to put a liquid crystal display (LCD) panel in front of a conventional optical see-through display in order to mask any rays coming from outside. This is not enough, however, since the LCD panel is so close to the eyes that a pattern on the panel is out of focus when a user sees outside objects. We conquered this problem by locating two convex lenses with the same focal length in front and behind the LCD panel. This optical system makes a telescope of one magnification. Finally, we use an erecting prism to erect inverted outside scenery. With this optical system, a viewer can simultaneously observe outside scenery and a pattern on the LCD panel clearly. Mutual occlusion of real and virtual worlds can truly be presented optically by opening and closing pixels on the LCD panel where real objects should appear and disappear respectively.




Figure 1: A prototype display.




Figure 2: Four image patterns.




Figure 3: A compact optics design of the display.

Our approach has three advantages:

- The display does not affect the real environment nor require special environmental settings at all. Therefore, the display is available in any situation, including outdoor applications.
- The display can mask rays from outside scenery in any situation. Therefore, color fidelity of virtual images is far superior to conventional optical see-through displays.
- The display can be used with most existing optical see-through displays.

Figure 1 shows our first prototype display. Field of view, horizontal and vertical viewpoint offsets, eye-relief, and transparency of open/closed pixels in the display are 35 degrees, 29 cm, 3.5 cm, 6 cm, 18%, and 2%, respectively. By using a PC (SGI VWS540) and a conventional optical see-through display (Shimadzu STV-E), four images (shown in Figure 2) were rendered. A white real box is modeled in the system and tracked by a 3D magnetic tracker (Polhemus Fastrak). Conventional optical see-through displays present a combined image like Figure 2(A). As shown in Figure 2(B), with depth information from a real object, virtual objects can be covered with the real one, but they remain semitransparent ghosts. Our display can mask the real scene as shown in Figure 2(C), and finally, as shown in Figure 2(D), it can present a mixed world with correct occlusion and high color fidelity.

CONCLUSION AND FUTURE WORK

Through these empirical studies, we confirmed that our novel display design solved the occlusion problem of conventional optical see-through displays. We have already designed and started to build compact optics, as shown in Figure 3. This lighter and smaller optical system allows a user to see a mixed-reality environment without viewpoint offset, and it will open the door to the practical mixed-reality world we will see in the near future.

www.crl.go.jp/ts/ts221/kiyo

References

1. Azuma, R.T. (1997). A survey of augmented reality. *Presence*, 6 (4), 355-385.
2. Inami, M. et al. (1999). Head-mounted projector, *ACM SIGGRAPH 99 Conference Abstracts and Applications*, 245.
3. Noda, S. et al. (1999). An optical see-through mixed reality display with a realtime rangefinder and an active pattern light source. *Trans. Of the Virtual Reality Society of Japan*, 4 (4), 665-670, (in Japanese).


BLUI (Body Language User Interface) has been developed to allow an artist to paint in the CAVE virtual environment. It is used to produce art. Virtual paint hangs in space, flowing from a wand in the hand of the artist. It is not bound to the accidental intersection of a painting instrument and a surface. BLUI currently recognizes “yes” and “no” gestures and is responsive to the speed and angle of the drawing instrument. Current efforts focus on more inclusive gesture recognition.

Tubes, lines, and multiple tubes are responsive to the angle and speed of the drawing instrument. Color is a function of speed filtered over time. Size is based on the angle between the instrument and the path swept out by its tip. Surfaces are generated by choosing any object and moving it. Copies of the selected object represented by a line are dropped based on simple distance criteria and a surface is lofted between the lines. This surface is edited by moving the whole or by displacing vertexes. The final object is written out as a Maya file. The object is imported into Maya and texture-mapped to appear like oil paint. The Maya scene contains a camera at its center that rotates one degree per frame and renders out a one-degree vertical slice 200 pixels

wide by 7200 pixels high. A python script calls on Image Alchemy Tools to glue 360 pieces together into a panoramic image and rotate the whole. Finally, Image Alchemy is used to create a file with suitable scaling and gamma correction in RTL format for printing on an HP DesignJet 3500. RTL is HP’s streaming printer control language. This file is sent to the printer, where it emerges as a 7200 pixel wide by 72,000 pixel high print four feet wide by 40 feet high, which is mounted on the inside of a 13-foot diameter cylindrical display so the immersive experience from the BLUI environment can be seen in a high-quality rendering.

This presentation includes:

1. A very-large-scale rendering of a scene painted in BLUI, depicting a panoramic view. The scene conveys a sense of the immersive quality of the BLUI experience.
2. Videotapes of an artist at work with BLUI that show gesture recognition and sensitivity to body language.
3. BLUI running on an ImmersaDesk.



Maya rendering of virtual paint created in BLUI.


The Panoscope 360° is a panoramic viewer capable of displaying a full 360° cylindrical image from a single video or data channel. A donut-shaped rear projection screen is placed horizontally above a conic reflector. From within the apparatus, visitors see a virtual cylindrical image obtained by reflexion.

Because it usually requires several film or video channels, production and presentation of 360° panoramic movies has been a complex and costly operation. The proposed apparatus, and its related imaging process, aims at simplifying both production and presentation of panoramic video programs. Such single-channel panoramic imaging technology will likely foster the emergence of a new type of broadly distributed content where visitors experience visual immersion.

The Panoscope 360° is particularly suited for interactive content where visitors are invited to navigate spaces and manipulate objects within them. It can be seen as a practical alternative to head-mounted displays by allowing several users to experience the same interactive content without having to wear special devices.

The Panoscope 360° consists of a standard XGA video projector, a computer, and a microphone. The special optical components include a conic reflector, a donut-shaped rear projection screen, and a conical reflector. The donut-shaped image used as source can be recorded with standard professional video cameras and panoramic lenses such as those from Cyclovision Technology (www.cyclovision.com) and Be-Here Corporation (www.behere.com). The video editing process is standard. Video content is stored on the computer's hard disc and manipulated with a microphone and voice-recognition software.

The demo will be featured in Space by Number, an interactive video program where visitors are invited to explore a territory by calling numbers displayed in the image.



*Panoscope 360° (patent pending).
Concept and design: Luc Courchesne.
Illustration: Etienne Desautels.*

For denizens of 3D space, visualization of figures in the fourth dimension is difficult. Dimensional analogy becomes an important tool; patiently pursued, it will gradually open a portal into this mysterious and beautiful fourth dimension.

Observing that a line segment is bounded by points, a polygon, by line segments, and a polyhedron, by polygons, by analogy we deduce that a 4-polytope is bounded by polyhedra. Just as a 3-space admits an infinite number of different planes, so also a 4-space admits an infinite number of different 3-spaces. From a mathematical standpoint it is convenient to represent an $(n-1)$ -space within an n -space in terms of a unit normal vector and a distance from the origin. For instance, given the unit vector $\{0,0,0,1\}$ in a 4-space, an infinite number of parallel 3-spaces exist, perpendicular to this vector.

To visualize 4-polytopes, we may project them into planes or into 3-spaces, with due regard to the distortion imparted by foreshortening. We may also obtain solid sections of a 4-polytope. There are sixteen regular 4-polytopes: six are convex, and 10 are starry. It is common to make projections of the six convex figures, such as the tesseract, or the 120-cell (bounded by Platonic dodecahedra), but rare to find constructions for the 10-star polytopes, much less projections or sections.

Why have these lovely figures been so neglected? They are much more difficult to construct, or to project, or to section, than their convex brethren. Star polyhedra in 3-space are difficult in their own right: of the four regular star polyhedra, two were discovered by Kepler, but centuries passed before Poinot found the other two.

I used *Mathematica* to construct the regular polytopes in four dimensions and developed methods to obtain solid sections. The star polytopes present special problems; their bounding polyhedra interpenetrated, and a single section might consist of disjoint polyhedra, or polyhedra with holes.

It became apparent that an animated series of equally spaced solid sections would produce a “morphing” of star polyhedra. Since the polygons in a section often numbered in the hundreds, or even thousands, and often intersected one another, I exported the sections into POV-Ray format, triangulating each polygon. A separate `#include` file was written for each section, 500 sections altogether.



A solid section of the star polytope $(5, 5/2, 3)$.

BACKGROUND

Ivan Sutherland, in 1970, dreamed of Virtual Environments (VEs) as a “window” to a virtual world that will ultimately “feel real.”⁶ But to feel real, a VE need not look real. The notion of presence in VEs conceptualises the feeling of “being there” and is defined as the degree to which the users think they are somewhere other than where they physically are while experiencing a computer generated simulation.⁵ Presence, although hard to define, offers a common attribute for evaluation of diverse VE design strategies. According to the recent virtual reality review by Frederick Brooks, one of the “hot, open challenges” is to “measure the illusion of presence and its operational effectiveness.”¹ However, to relate presence, explained as a psychological state of consciousness, to a specific measure, allowing a quantitative assessment, introduces a hard problem. Investigating a subjective notion, commonly assessed by post-experience self-report or magnitude estimation methods that are non-intuitive for the user, does not provide by itself concrete information for the man-machine quality of communication, without task-performance measures. Moreover, relevant research has not yet established a clear, positive correlation between the notion of presence and task performance, associated with real-world applications.

EXPERIMENTS AND RESULTS

The methodology presented here focuses on real vs. simulated virtual worlds in order to obtain users’ response data on their level of presence as well as task performance, acquired uniformly under three distinct conditions. A better understanding of how to design and implement VE systems will result from discovering links between presence and human performance, as well as incorporating real-world data into VE applications.

A set of pilot studies investigated the importance of visual and auditory stimuli on memory and the sense of presence. Three different sets of 18 subjects completed the same memory task, the first set in reality, the second using non-realistically rendered 3D with audio on a typical display, and the third under audio-only conditions. Users attended a 15-minute seminar and completed two questionnaires in each condition. The first one evaluated their memory recall of the seminar and space as well as their confidence level and memory awareness state, for every question. The second one determined the users’ level of presence. Conclusions were based on analysis of variance of the resulting data:³

- Presence, assessed by a Likert-scale questionnaire, correlated positively with task performance in most (but not all) conditions. Presence and task performance were higher (and statistically different) in the “real-world” condition.

- Presence for the 3D vs. audio condition was not statistically different overall. This is attributed to the non-realism of the scene as well as to the mental-imagery cognitive process connected with the audio sensory input. Task performance was higher for the 3D condition vs. audio but only for the aspects of the task that require the visual stimuli to be there.
- Presence proved to be different than “engagement” or “enjoyment.” This result validates the assessment method for presence.

The follow-up set of experiments employs a visual search/memory task of 3D primitive objects (spheres, boxes, etc.). The task is conducted in a real room but also in simulated 3D “rooms” with a goal of using different types of display technology. In general, the degree of “immersion,” by itself, does not provide any information about the degree of presence or its relation to task performance, as it only offers a description of VR technology. Thus, presence assessments as well as objective measures of users’ task performance, combined with cognition-related responses, are acquired in each condition, including the real-world one. Since the optic array from a picture and the optic array from the real world can provide the same information or impression without providing exactly the same stimulation,² optimal distribution of computational power can be identified.

A high degree of presence would enable users to behave in a VE in a manner similar to their behaviour in everyday reality.⁴ This is significant when designing simulation and training applications, especially toward transfer of training. Therefore, the relation of presence and task performance is explored, here, in order to identify appropriate uses of available technology and classify suitable VE applications.

ACKNOWLEDGEMENTS

This research is funded by the Hewlett-Packard Laboratories External Research Program. Thank you to Mel Slater and Martin Usoh for provision of the presence questionnaire.

References

1. Brooks, F.P. Jr. (1999). What’s real about virtual reality. *IEEE Computer Graphics and Applications*, Nov-Dec. 1999.
2. Gibson, J. (1971). The information available in pictures. *Leonardo*, 4, 27-35.
3. Mania, K., Chalmers, A. (2000). A user-centred methodology for investigating presence and task performance. *3rd International Workshop on Presence*, ISBN 9-386-1571-X.
4. Pausch, R., Proffitt, D., & Williams, G. (1997). Quantifying immersion in virtual reality. *Proc. SIGGRAPH 97*.
5. Slater, M., Usoh, M., Linakis, V., & Kooper, R. (1996). Immersion, presence and performance in virtual environments: An experiment with three-dimensional chess. *Proc. VRST 1996*.
6. Sutherland, I.E. (1970). Computer displays, *Scientific American*, 222 (6).

Recent work in computer-animated characters^{1,2} has demonstrated interest in splitting their behaviour into two parts: one that computes it and one that expresses it. The modularity thus fed into the system changes the assumptions about the system as a whole. We focus on personality modeling, and especially its relationship with the interaction process. Here, we present our work in progress in the field of communication (visual and textual) and synthetic personality building, in the context of the “2ème monde,” the Canal+'s French speaking virtual community.³

RESEARCH CONTEXT

The “2ème monde” is built on the Blaxxun virtual communities technology.⁴ It is a chat- and 3D-based virtual community, so we are considering two modalities of interaction: visual (through an avatar and body expressions) and textual (like in any chat-room). The body expressions can be standard or user defined, by using a dedicated avatar modeler.⁵



Avatar waving goodbye.

Our aim is to populate cyberspace with artificial and semi-artificial entities (explained below). To have these entities accepted by humans, we think that they have to respect some of the standards of human interaction. Firstly, they have to consider messages as interpretable — messages may not be what they first look like. Secondly, they must exhibit some soundness as well as some politeness or social awareness.

AGENTS, PERSONALITY, AND INTERACTION

Two kinds of mind can inhabit an avatar: humans (natural agents) or computer programs (artificial agents). Nevertheless, we also consider “composite entities,” semi-artificial entities whose avatar is sometimes controlled by a human, and sometimes by a program trying to mimic the behaviour of the human user. The personality module we are working

on will be included in the next versions of Avatar Studio⁵ in order to build the personality of an avatar at the same time as its body. We express the personality through variables derived from psychological research, such as eagerness of domination or tendency to friendliness. While most studies of Interactions between agents consider that interacting is mainly reasoning, we think that we also have to take into account the mechanisms of interpretation and expression. We consider a message as a holder of a potential meaning (rather than of a simple meaning), which depends on other parameters than the message itself: the relationship between the two interacting agents (social awareness), their personality and current “state of mind,” and their previous exchanges (soundness). For example, if a friend of yours says “I’d like to go to the movies,” it looks like a suggestion, and can be interpreted as one. But if a mother says to her son “I’d like you to clean your room,” it still looks like a suggestion but may be interpreted as an order!

CONCLUSION AND EXPECTATIONS

By attaching human-like abilities to bots, we are bringing them closer to humans, thus easing their interaction.⁶ The 3D part of the interaction holds the specificity of the 2ème monde, compared to usual chat-rooms. It is important that our bots take advantage of this modality for two main reasons: First, because multi-modal communication greatly eases the interaction process. Second, because visual interaction is mostly unused by the human users (because of the design of the interface, but also because the users themselves are not really aware of the capabilities of this medium). So, as a side effect of our work, we would like to enhance visual interactions between the users. By exhibiting a visual behaviour, we believe that our bots will encourage the use of this modality by humans.

References

1. OZ Project publications, available at: www.cs.cmu.edu/afs/cs/project/oz/web/papers.html
2. MITsoftware agents group publications, available at: lcs.www.media.mit.edu/groups/agents/
3. Le deuxième monde: virtuel.cplus.fr/
4. Blaxxun: www.blaxxun.com
5. Avatar Studio: www.avatarstudio.com
6. *Proceedings of IWA'I'99 (International Workshop on Affect in Interaction)*, Sienna. Available at: gaiva.inesc.pt/i3ws/accepted-papers.html

We present a novel approach to segmenting and modeling the 3D anatomy of the Visible Human data⁴, which leads to a nearly photorealistic, interactively explorable model. The well-known volume-rendering approach fails to display complex internal structures properly. In addition, organs cannot be addressed, thus making removal or exclusive display of an organ impossible. So segmentation (exact determination of the surface location of an organ) is mandatory for its realistic display.

MATERIAL AND METHODS

Data: We used a subvolume of the photographic cross-sectional images of the male Visible Human together with registered computer tomography data from the fresh and frozen cadaver to create a model of the torso containing the inner organs.

Segmentation: We found that objects in the photographic image volume are characterized by clusters in RGB space that have an ellipsoid shape. The volume was segmented with an interactive segmentation tool² that converts the resulting clusters into the mathematical form of an ellipsoid to speed up the rendering process.

Graphic modeling: In order to arrive at a complete model, we developed a modeling tool, which allows inclusion of non-segmentable tube-like objects like nerves and small blood vessels into the model.

Knowledge modeling: Spatial representation of anatomy is connected to a knowledge base that contains object descriptions within the structure of a semantic network¹. Different networks exist for different views (for example, systemic or regional anatomy). Within the views, the anatomical constituents are linked by relations like "part of" or "branching from."



Figure 1: The model allows arbitrary views into the complex anatomy using cut planes and the ability to remove or add objects.




Figure 2: Different viewing modes such as X-ray imaging may be chosen from any direction for any part of the image.

Visualization: We have developed a 3D visualization algorithm based on the ray casting paradigm that renders surfaces from volume data. In contrast to volume rendering, the rays stop at the first encountered object. Surface texture and surface inclination are calculated from the RGB tuples at the segmented border line. Decisive quality improvement is achieved by determining the object surfaces with a spatial resolution higher than the resolution of the original voxels. While such subvoxel resolution may be achieved easily by sampling with high rates in volume rendering, determination of adjacent surfaces in segmented volume data is a major problem that has been solved only recently⁵.

RESULTS

We have built a model of the organ systems within the torso of the male Visible Human. It contains more than 400 3D space-filling anatomical constituents, 300 graphical objects consisting of about 500,000 polygons, and more than 2,000 relations among the objects. The model is novel because:

- (Figure 1) The visual impression is one of unsurpassed realism, due to the exact segmentation and the visualization method developed.
- The model is space-filling (each voxel is labeled as an element of a 3D object).
- The integrated formal organization of pictorial and descriptive information allows virtually unlimited ways to use the model.
- The model may be used interactively for teaching and simulation of radiologic and interventional procedures³. It is the first to offer sufficient detail and comprehensiveness for professional use (Figure 2).

It will be available in summer 2000 as VOXEL-MAN 3D-Navigator: Inner Organs - Regional, Systemic and Radiological Anatomy at Springer Electronic Media.

References

1. Hoehne, K.H., Pflesser, B., Pommert, A., Riemer, M., Schiemann, T., Schubert, R. & Tiede, U. (1995). A new representation of knowledge concerning human anatomy and function. *Nature Med.*, 1 (6),506-511.
2. Schiemann, T., Tiede, U., & Hoehne, K.H. (1997). Segmentation of the Visible Human for high quality volume based visualization. *Med. Image Anal.*, 1 (4),263-271.
3. Schubert, R., Pflesser, B., Pommert, A., Priesmeyer, K., Riemer, M., Schiemann, T., Tiede, U., Steiner, P., & Hoehne, K.H. (1999). Interactive volume visualization using "intelligent movies." In Westwood, J.D., Hoffman, H.M., Robb, R.A., & Stredney, D. Eds., *Medicine meets virtual reality, Proc. MMVR '99*, 62, of *Health Technology and Informatics*, 321-327. Amsterdam: IOS Press.
4. Spitzer, V., Ackerman, M.J., Scherzinger, A.L., & David Whitlock. (1996). The visible human male: A technical report. *J. Am. Med. Inf. Ass.*, 3 (2), 118-130.
5. Tiede, U., Schiemann, T., & Hoehne, K.H. (1998). High quality rendering of attributed-volume data. In Ebert, D., Hagen, H., & Rushmeier, H. Eds., *Proc. IEEE Visualization '98*, 255-262.

Various methods have been proposed for the modeling and animation of human hair, including methods of modeling hair strands individually,^{1,2} trigonal prisms with 2D hair distribution map,³ and more recently, volumetric visualization techniques.⁴ These approaches concentrate on modeling hair accurately, often requiring very high computing power, thus making them more suitable for off-line graphics system.

We propose a novel framework that is more suitable for real-time applications. The main idea is to model and animate the hair in strips, using B-Splines, texture-maps, and transparency, taking advantage of common 3D hardware acceleration.

PROPOSED FRAMEWORK

Hair strands can be grouped into strips, which are modeled geometrically by NURBS surfaces. Surface patches may have different shapes and sizes representing hair strands on different parts of a head model. We provide functions and a user interface to model and attach the hair strands to the head model.

Texture maps of the hair images can be applied on either one or both sides of each surface patch. The alpha map defines transparency and creates an illusion of complex geometry to the otherwise “rectangular” surfaces. It adds realism to the final result (Figure 1).

We used the advantages of parametric surfaces. One of them is that the level of detail is scalable by the polygon retessellation. The MIP-Mapping can be applied for texture mapping.⁵ This allows for

real-time performance. Different combinations of geometry and texture maps may be defined for each hair strip.

Key-frame animation (Figure 2) is achieved by displacing the control points of hair strips. Existing animation techniques can be applied such as physically based modeling of clothes.⁶ We are exploring further in this direction. The framework is also applicable to modeling and animation of leaves or grass.

IMPLEMENTATION

The Java3D 1.1.3 API is used for the implementation, with each hair strip defined as one object containing methods to retessellate and deform its geometry. The listeners wake up on events to trigger the methods. Finally, the object-oriented model improves interactivity.

References

1. Anjyo, K., Usami, Y., & Kurihara, T. (1992). A simple method for extracting the natural beauty of hair. *Proc. SIGGRAPH 92*, 111-120.
2. Daldegan, A., Kurihara, T., Magnenat Thalmann, N., & Thalmann, D. (1993). An integrated system for modeling, animating and rendering hair. *Eurographics 93*, 211-221.
3. Chen, L. H., Saeyor, S., Dohi, H., & Ishizuka, M. (1999). A system of 3D hair style synthesis based on the wisp model. *The Visual Computer (99)*, 15 (4), 159-170.
4. Kong, W. & Nakajima, M. (1999). Visible volume buffer for efficient hair expression and shadow generation. *IEEE Computer Animation (99)*, 58-65.
5. Williams, L. Pyramidal parametrics. *Proc. SIGGRAPH 83*, 1-11.
6. Baraff, D. & Witkin, A. (1998). Large steps in cloth simulation, *Proc. SIGGRAPH 98*, 43-54.




Figure 1: Modeling Hair



Figure 2: Snapshots of four frames from the key-framing animation.

BACKGROUND

Aerodynamic effects in games and films are difficult to generate and time consuming to produce due to the inherently non-linear and complex nature of the Navier-Stokes equations. Although faster and stable solutions are appearing,¹ these invariably involve solving a form of the Navier-Stokes equations over a grid. Although they are fast to implement, grids require that solid object-boundary conditions must be applied at grid points, which requires high resolution for complex or curved objects. In addition, moving objects must be re-gridded. Both procedures are computationally expensive.

The MathEngine fluid-dynamics library overcomes the gridding problem by using a combination of classical steady-state solutions for large object flows and a Navier-Stokes-based particle system for the unsteady part of the flow. The steady-state solution is found using an irrotational potential flow model for an object in uniform flow. The unsteady flow is produced by tracking the movement of vorticity in the form of spherical vortex particles. These are originally placed into the flow according to empirical models. The result is a real-time method for fluid-flow effects.

EXAMPLE

As a ball travels through air, it disturbs the airflow in front of it and creates a wake behind it in the form of turbulent eddies. We can reproduce such a flow system by simply specifying parameters for a sphere, sphere

wake, and main flow. By altering the position and velocity state of the ball, the user can interact with the flow in real time. The advantage of the vorticity-based particle system over traditional particle systems is that complex interactive wake behaviour can be reproduced. If another ball is placed just behind the original ball it will affect the original wake as well. More complex examples, such as the trailing vortex behind a wing airfoil or flat plane deflecting a flame jet, can be modeled in the same manner.


Mass-less test particles can be introduced to evaluate the flow velocity at any position. This lends itself well to a simple rendering system where textured billboards are convected along by the flow as in Figure 1. Also, test points can also be used to generate approximations to aerodynamic forces from velocities. The example given above will create unsteady aerodynamic force on the second ball.

FUTURE

Such a system can allow for far more complex shapes in the flow field than grid-based solutions. Flows around geometry such as racing vehicles and inside office buildings are possible examples of applications in the near future.

Reference

1. Stam, J. Stable Fluids, *Proc. SIGGRAPH 99*, 121-128.



The uniform flow in the form of billboard test particles is seen approaching from the left. The turbulent wake behind the object on the right is clearly visible.

Traditionally, high-tech graphics emphasize the most realistic, lifelike images that can be created. This project is one of many that have explored a new direction in computer graphics: non-photorealistic rendering. Many attempts in this direction have dealt with filters applied over 2D images, but we have opted to build a rendering engine that creates a virtual world comprised of 3D brush strokes. This gives the entire virtual environment an aesthetic of having been created by a human painter while still maintaining an illusion of solidity and continuity.

Starting with standard Lightwave/OBJ input, the system processes data through an “Artist’s Interface.” This allows the user to decide what types of brush strokes should be applied to which objects in a scene. Variables such as stroke length, angle, color, texture, and size can all be defined to the polygon level. Once the user has defined the specifics for the objects in the 3D painting, these objects are sent through a tessellation algorithm, broken into individual brush strokes, and configured with the proper brush settings.

After this preliminary configuration, the rendering engine computes all of the rotations and translations needed to allow movement in the scene. Once the appropriate matrix arithmetic has been completed, the engine then plots the actual brush strokes to the screen with the proper size, color, angle, and texture. At this stage, our system also applies various real-time filters to the brush strokes to generate the effect of dynamic changes in a 3D world. These filters, for example, can give rivers the effect of flowing, oceans the effect of glittering, and leaves the effect of rustling. We also have a filter that allows objects in the scene to be modified in real time based on sound; objects can be

defined to change their brush strokes’ angle, length, texture, and color based on the intensity and frequency of sounds acquired with a microphone.


The end product is a virtual environment that is fully navigable, giving the user a feel of being in a 3D piece of art. We have been able to create interactive paintings that emulate the work of famous artists (for example, Van Gogh and Monet). A reasonable painting can run as fast as 10 frames per second on an entry-level, 3D-accelerated desktop system. Simple paintings can reach speeds as fast as 30 fps including dynamic effects. The system also contains an option to dump raw painting data to disk, allowing creation of videos that run at 60fps. The system is scaleable to complexities ranging from single-stroke to million-stroke scenes.

This project was programmed by a team of eight students in the University of Illinois SIGGRAPH Chapter. Over 2,500 hours were invested in the final version of the project. The project took second place in undergraduate research at the university’s Engineering Open House exhibition.

This is an approach much similar to that described by Barbara J. Meier in “Painterly Rendering for Animation.”¹

Reference

1. *Painterly Rendering for Animation*. Barbara J. Meier, Walt Disney Feature Animation. <http://www.acm.org/pubs/articles/proceedings/graph/237170/p477-meier/p477-meier.pdf>



Reconstruction of the World's First Panoramic Immersive Virtual Environment Using Contemporary VR Technology

Mark Wright
Gordon Watson
Roy Middleton
University of Edinburgh
mark.wright@ed.ac.uk

One of the most challenging problems in computer graphics is to give people the impression that they are somewhere else. Ever since Ivan Sutherland proposed the creation of an “ultimate display,”¹ the graphics community has made realistic rendering of environments a priority. Virtual reality (VR)² introduced² the concept of immersion, where we no longer feel we are looking at an image of an environment, but that we are in that environment.

More recently, image-based rendering³ has meant that photo-realistic representations of scenes can be rendered. Finally, semi-immersive, large-screen projection systems⁴ have allowed groups of people to be immersed simultaneously. It is surprising to discover that the root concepts behind these advances belong not to the 20th century, but to the end of the 18th century. The world's first panoramic, immersive virtual environment was devised and built in Edinburgh, Scotland over 200 years ago.

The inventor of these concepts was the painter Robert Barker. As he considered how he might capture a vista of Edinburgh in its entirety, he devised a cylindrical perspective with which he constructed a huge painting (Figure 1). A patent was granted in 1787,⁵ which describes how the painting should be viewed. The “panorama” was hung in a large circular building called a rotunda (Figure 2), with the painting wrapped around the inside wall like the label on a soup tin. People viewed the panorama from a central platform. These panoramic immersive virtual environments caused a sensation and were very popular in Europe and America for the duration of the 19th century.

Unfortunately, none of the full sized panoramas has survived, but there are some small panoramas measuring 0.3 meters high by 2.5 meters long. We have recreated an immersive panoramic virtual environment from one such painting. The original panorama was photographed in sections using a medium-format camera, and the film was then digitized. A complete panorama was created from these separate photographs. The digital panorama was re-mastered by an artist, to restore damaged areas. We then modeled the rotunda by projecting the panorama onto a cylinder. The panorama was viewed using a tracked head-mounted display. This is the first time the original panorama has been seen as a panoramic immersive virtual environment for over two centuries.

References

1. Sutherland, I. (1965). The ultimate display. *Proceedings IFIP Congress*, 506-508.
2. Sutherland, I. (1968). A head-mounted three dimensional display. *Proceedings AFIPS Fall Joint Computer Conference*, 757-764.
3. Chen, S.E. & Williams, L. (1993). View interpolation for image synthesis. *Proc. SIGGRAPH 93*, 279-288.
4. Cruz-Neira, C., Sandin, D.J., DeFanti, T.A., Kenyon, R., & Hart, J.C. (1992, June). The CAVE, audio visual experience automatic virtual environment. *Communications of the ACM*.
5. Barker, R. (1796). *Repertory of arts and manufactures* 4, 165-167. (Patent granted 19 June 1787).




Figure 1: Original panorama of Edinburgh. The panorama was used to create an immersive virtual environment using a technique patented by Robert Barker in 1787.




Figure 2: A rotunda. The rotunda was a large circular building in which the panorama was hung around the wall. People could view the painting from a central platform. Viewers could not see the top or bottom of the painting, which created the illusion of immersion in a virtual environment. 3D objects were added in the foreground to intensify the illusion.

Digitizing visual data has had a dual impact. First, it has enabled fast and efficient storage, transfer, and processing. Secondly, duplication and manipulation of data have also become very easy and undetectable. This can lead to violation of copyrights. One solution to this problem is to embed a secondary signal (watermark) into the visual data in an unobtrusive and robust manner.

We have developed a novel algorithm, based on the spread-spectrum communication technique, that is invisible and robust. "Invisible" means that the 2D volume-rendered image of this watermarked volume is perceptually indistinguishable from the volume-rendered image of the original volume. "Robust" watermarking implies that the watermark is resistant to intentional or unintentional attacks.

THE VOLUME WATERMARKING TECHNIQUE

Assume that volume V that needs to be watermarked:

- (a) A block-based 3D discrete cosine transform (DCT) is performed on the volume V .
- (b) To embed the watermark bits, the bits are first replicated by a large spread factor and then modulated with a pseudo-random-noise sequence.
- (c) This watermark sequence, which can be considered to be a volume, is also 3D DCT-transformed.
- (d) For every DCTblock of (a) and the corresponding DCTblock of (b), the coefficients are added to form a watermarked block that constitutes the watermarked volume in the frequency domain. The 3D inverse DCT is performed to obtain the watermarked volume.

Since a pseudo-noise sequence is used for modulation, the watermark is noise-like, which ensures that it is difficult to detect, locate, and manipulate without compromising the corresponding volume-rendered image quality.

WATERMARK DETECTION

To detect the existence of the watermark, the DCT-transformed original volume data are subtracted from the DCT-transformed watermarked volume data to obtain the residual volume-data DCTcoefficients. The 3D inverse DCT is performed on this residual data to obtain the residual watermark sequence. This is then analyzed by correlating it with the same pseudo-noise sequence that was used in the embedding phase.

RESULTS AND DISCUSSION

The quality of the watermarked image appears unaltered as seen in Figure 1. The images on the left side are volume-rendered images of the original skull volume data set. A 170-byte watermark was embedded into this volume dataset using the proposed technique. The images on the right are the volume-rendered images of the watermarked skull volume data set. We have done extensive experiments that indicate that the perceptual quality of the volume-rendered images of the watermarked volume datasets is largely unchanged.

We are still refining our technique and doing further experiments to test the robustness against various attacks.




Figure 1: Volume rendered images of the original volume.




Figure 2: Volume rendered images of the watermarked volume.

As we begin to create synthetic characters (avatars), it is important to model subtle non-verbal behaviors. One of the most important behaviors is the look and the motion of the avatar's eyes. We present behavior models of eye-gaze patterns in the context of real-time verbal communication and apply these eye-gaze models to simulate eye movements in a computer-generated avatar. We have relied heavily on a broad body of literature to develop this eye gaze model. We also report the results of an experiment conducted to assess whether our eye-gaze model induces changes in the eye-gaze behavior of an individual who is conversing with an avatar.

The stochastic eye-gaze models we developed are summarized in the hierarchical state machine diagrams shown in Figure 1. Transitions between major states are triggered by a change in which participant is talking. Built into each major state are lower-level states indicating where the avatar is gazing and the state of the other participant. The decision to transition the avatar's eye gaze (that is, whether it looks at or away from the other) is triggered primarily by duration in the current sub-state. As one might expect, each of these timing values varies both within a particular person and across persons and pairs of persons.

A first question that arises is whether users interacting with an avatar will act in a way that resembles interaction with a human. To investigate this question, we conducted the following experiment (illustrated in

Figures 2 and 3). A subject sits in front of a monitor displaying an avatar and converses with an experimenter. A second experimenter tracks the subject's eye gaze. The microphone input of both conversant and subject, along with the subject's eye gaze, is fed into the model, and drives the eye gaze of the avatar. The subject was shown three three-minute segments: a blank screen, an avatar that always looked at the subject, and one with a functioning eye-gaze model. A total of 20 subjects participated in the study. Not surprisingly, the subjects looked at the monitor more when an Avatar was present; 64% vs. 38% of the time. The subjects also looked at the avatar more often while listening when the gaze model was active vs. when there was no gaze model. The increase was from 67% to 74% of the time. This difference was surprisingly consistent across subjects. T-Test results were 0.1, indicating a 90% chance that a significant difference was detected. There was also a difference in the ratio of time spent looking while listening vs. looking while speaking. The ratios were 1.2 without the eye-gaze model and 1.3 with the model, with a T-test of .17, which although not conclusive is significant. It is of interest that we found these results even though only one subject had consciously noted any change in the eye-gaze pattern.

As we begin to build realistic virtual human characters, it will be increasingly important to have them behave in natural ways. This sketch focuses on one aspect of that behavior.




Figure 1




Figure 2



Figure 3

Simulation is an important method for animating virtual environments, but its computational cost limits the size and complexity of highly dynamic worlds. We are developing new algorithms for simulation culling that are applicable to objects that move over large distances and exhibit group behaviors, such as the approach described here for cars in a city. Through culling, we can scale simulation to very large environments while paying little for motion out of view.

In previous work,¹ we describe simulation culling for independent objects with small spatial ranges, which precludes cars in a city. For simulation level-of-detail, other authors^{2,3,4} have used cheap proxy simulations that still work for every object on every frame and compute full group behaviors on the approximated motion. For dependent, wide-ranging objects, we improve upon previous results in two ways:

- Only objects in view are fully simulated on every frame, while objects out of view are updated at a lower frequency with much lower cost.
- Probabilistic models are used to capture the overall effect of group behaviors on objects out of view without explicitly tracking interactions.

We experiment with automatically generated cities populated by cars that roam the streets making random turns at each intersection. The cars obey the rules of stop-sign intersections and avoid collisions by adjusting their speed. On each frame, the potentially visible roads are made known to the car simulation, which then maintains a list of potentially visible cars. Every other car makes a guarantee, marking roads that it may be on up to some future expiration time. The guarantees are computed by doing breadth-first search on the road network to locate roads the car may be able to reach in the future. Bounds on maximum velocities allow us to say how long it will take to reach any given road, and hence compute expiration times.

On each frame we perform the following operations:


- Make new guarantees for all cars with expired guarantees.
- Determine the location of each car whose guarantee includes potentially visible roads, and, if visible, add the car to the potentially visible set or make a new guarantee.
- Simulate the potentially visible cars as normal.
- Remove any cars that are no longer visible from the potentially visible set and make guarantees for them.

To position a car when its guarantee overlaps the view, we associate a probability distribution with the travel time for each road and do random walk with sampled travel times. The distributions are based on statistics gathered from an off-line simulation of all the cars, so they directly encode the cumulative effects of group behaviors. Also associated with each road are various approximation functions that allow us to set the state of the car after it has been out of view.

Without culling, we can simulate at most a few hundred cars on a PC. With culling, we can scale up to thousands of cars. The performance of a culling algorithm is best defined by its efficiency, which is the ratio of work done for systems in view to total work done. We see efficiencies ranging from 86% to 96% in experiments performed with a constant-sized city and a variable number of cars. In experiments using larger cities with up to 2,400 cars, efficiencies fall to 35% because we track all the objects that may ever be seen in the world. In future work, we expect to address this by tracking only those objects the viewer has seen recently and adding and deleting other objects as needed.

References

1. Cheney, S. Ichnowski, J. & Forsyth, D. (1999). Dynamics modeling and culling. *IEEE Computer Graphics and Applications*, 19 (2), 79-87.
2. Grzeszczuk, R., Terzopoulos, D., & Hinton, G. (1998). NeuroAnimator: Fast neural network emulation and control of physics-based models. *Computer Graphics, Proc. of SIGGRAPH 98*, 9-20.
3. Carlson, D.A. & Hodgins, J.K. (1997). Simulation levels of detail for real-time animation. *Graphics Interface '97*, 1-8.
4. Yu, Q. & Terzopoulos, D. (1999). Synthetic motion capture: Implementing an interactive virtual marine environment. *The Visual Computer*, 377-394.



Screen snapshots from a city simulation. The left image shows all of the city, populated by 100 cars with culling turned off. Red roads are potentially visible, and the shaded cone indicates the viewer's location. The middle image is a close up while culling is in progress. Only cars on or near the visible roads are being simulated. The green shades indicate how many non-visible cars may be on the other roads. The right image is the walkthrough view corresponding to the middle map. Note the car stopped at the intersection waiting for another to pass.

We address an open problem in image-based rendering: how to accurately render synthetic images of a real object under arbitrary lighting conditions. Our solution makes no assumptions about the object's shape or the presence or absence of shadowing, nor does it make assumptions about the nature of the BRDF.

Consider an object whose pose is fixed with respect to a camera. If we gathered an image illuminated by a single ray of light (single-ray image) and repeated this process for every light ray in the 4D light-ray space, we could then synthesize a images under arbitrary lighting by appropriate superposition of single ray images. But densely sampling the 4D light-ray space of an object is currently too data-intensive.

To bypass this, we have developed a dual to the Lumigraph (Light Field) method in which images are acquired as an isotropic point source (each covering a 2D set of light rays) and moved over a surrounding surface (a 2D set). Since each ray emanating from a point light source strikes the object at only one point, and since no two rays strike the object at the same point, each image provides a unique 2D sampling of the object's 4D light ray space. If the scene geometry, viewpoint and, point source position are known, we can decompose each of these images into a 2D set of single-ray images. The collection of images gathered as the light source moves over the surrounding surface yields all single-ray images for the full 4D light-ray space.

To establish the scene geometry, a novel reconstruction method has been developed based on a simple assumption: radiance along a light ray is constant. The depth λ of a scene point p is determined by moving an isotropic point source over two non-intersecting surfaces parameter-


ized as $s_1(\phi_1, \psi_1)$ and $s_2(\phi_2, \psi_2)$. For every light source position from the first surface $s_1(\phi_1, \psi_1)$, there is a source on $s_2(\phi_2, \psi_2)$ where the ray from p through $s_1(\phi_1, \psi_1)$ is identical to the ray from p through $s_2(\phi_2, \psi_2)$; see Figure 1. We can express this correspondence as a change of coordinates $\phi_2(\phi_1, \psi_1; \lambda)$ and $\psi_2(\phi_1, \psi_1; \lambda)$ parameterized by λ . The relation between the image irradiance under corresponding sources is used to form an objective function in λ .

$$\mathbf{O}(\lambda) = \int \int [i_2(\phi_2(\phi_1, \psi_1; \lambda), \psi_2(\phi_1, \psi_1; \lambda)) - d^2(p(\lambda))i_1(\phi_1, \psi_1)]^2 d\phi_1 d\psi_1.$$


where $i_1(\phi_1, \psi_1)$ and $i_2(\phi_2, \psi_2)$ respectively denote the measured image irradiance at pixel q for light sources at $s_1(\phi_1, \psi_1)$ and $s_2(\phi_2, \psi_2)$; $d^2(p)$ is defined below. Depth is recovered independently for each pixel by minimizing $\mathbf{O}(\lambda)$.

Once depth and $p(q)$ are available for each pixel, we have effectively determined which rays in the 4D light-ray space shed light onto each pixel. To render an image for arbitrary lighting (a positive function on the 4D light-ray space), the image irradiance for pixel q is determined by a weighted sum of the recorded irradiances for each source ray striking $p(q)$. Because the solid angles of a surface patch at $p(q)$ depend on the squared distance, recorded irradiance is scaled by $d^2(p(q)) = \frac{\|p - s_2\|^2}{\|p - s_1\|^2}$. It is straightforward to extend this to color images.

The methods have been tested using images acquired with a CCD camera, and a three-color LED source was accurately positioned using an ADEPT robot arm.



A 2D schematic of the reconstruction setup. A camera whose origin is at \mathbf{O} observes a scene point p that is illuminated by a point source moved over two surfaces.



Rendered image of a pear illuminated with two sources: a point source to the left and a 3 cm by 5 cm area source to the right.



Rendered image of a seashell illuminated by an isotropic point source off of the two surfaces.

In Paul Verhoeven's "The Hollow Man," we needed to show a chemical reaction that changes a human from visibility to invisibility and back again. One of the challenging aspects of this "transformation" effect involved rendering animating structures inside the human body. Our starting point was a highly detailed model of Kevin Bacon. The model consisted of an external skin and internal components such as skeleton, muscles, and organs. The transformation design required that we show a chemical reaction starting inside the body and flowing through the bloodstream until it finally revealed the solid form of a body. A model gave us the final outer surfaces for each element, but our main problem was creating the mass of tissues that lie inside bones, muscles, and organs.

At an early stage in the project, we were told that dissolving transparent objects was not the desired look, so simple techniques were not an option. The transformation also had to appear excruciatingly painful, which meant we would see a lot of writhing body parts with frenetic camera moves. Our solutions needed to work with rapidly deforming objects that could be viewed from any angle and distance.

Rendering the transformation sequence required a complex blending of medical visualization and cinematic-quality imagery. We had to take techniques from the medical arena to provide the solid volumes inside the body and combine them with more traditional surface rendering. From the beginning, we were aware that many volumetric datasets were available, but there are many limitations to taking these data and just plugging them in. Unless the data fit perfectly inside a model, they will be unusable. Another problem would have been the volume of data. Textures would no longer be manageable 2D elements, because they would become gigantic cubic pieces of data, requiring more costly 3D filtering.

Most MRI data, and especially sliced data, such as the Visible Human project, only represent the state of a volume at a frozen moment in time. If we were to show these volumes moving as a living being, we would have to extrapolate motion from these data. This would have meant taking every voxel and reverse mapping it through the distorted volumetric space it had traveled in order to apply a consistent color or density value. These calculations would have been very expensive and prone to large margins of error.

Faced with these challenges, we decided to leverage previous experience. We knew that with RenderMan shaders we could replicate an organic look using a combination of painting and procedural textures. Since these textures were procedural, they did not consume vast quantities of storage, and since they are cubic in nature, we only had to keep track of a cube deforming instead of an arbitrary shape such as a lung or kidney. We could also use the same lattice that was animating an organ to calculate the distortions in the cubic texture, thereby locking a texture to the object in the volumetric space. Once we had a "solid" object, we needed to be able to render it in various states of growth.

We decided to split our research into two paths using two renderers, RenderMan and a proprietary volume renderer. RenderMan was used for transforming bones, muscles, ligaments, and heart. To be able to use RenderMan for this task, we needed to give it the ability to ray trace and "ray march." We solved this problem by taking advantage of the DSO mechanism recently made available in RenderMan 3.8. We also

created a variety of DSOs that gave us access to data that allowed us to calculate the position of a voxel relative to other parts of the body. Using this data as animation properties, we could decide whether a voxel should be visible or invisible. This made it possible to take objects like bones or the heart from an empty space to a solid volume and back again without moving the surface geometry. By evaluating neighboring voxels we were able to create surfaces that could be lit and textured as if they were real geometry.

Vrender is the renderer we created for "solid" organs. These organs (for example, kidneys, liver and lungs) were to become visible by first showing the internal structure of arteries that fed into the cells of the organ. These arteries were created using Houdini SOPs to build a set of branches that fit inside the volume of an organ. Vrender gave us the ability to evaluate "blobby" surfaces derived from animating particle data. These blobs could then be volumetrically rendered with varying densities to create a sponge look for the lungs or a dense solid look for the liver. The surface was expanded out until it hit the outer geometry of the organ, where it blended with a RenderMan static organ look. This meant that both Vrender and RenderMan shaders needed to live in the same scene with similar lighting models, and shadows from each renderer had to be cast and received from the other.

To solve the problem of moving blood through the body, we essentially designed a system similar to a road map. We realized that we could not animate an object's visibility based purely on its position along a major axis. We needed to calculate the curved distances that veins and arteries take through the body. Our solution was to create a new curve class within Maya and generate a network of curves that would describe the distance from the heart to any part of the body. In the case of small veins, we rendered these curves as veins themselves. Another plug-in gave us the ability to create curves across multiple patch surfaces, so we could make veins that ran over the surface of muscles and bones with more control than an L-System. To edit the network, any curve could be attached, inserted, or detached, and the network would then recalculate its properties. At render time, these curves could be fed into RenderMan DSOs to either control animation or produce renderable curve primitives.

The transforming muscles needed to appear as bunches of fiber strands that kept growing until the space occupied by each muscle was filled. To make the problem more complex, the strands needed to grow inside muscles that were animating. Once the fibers were dense enough, they would also need to start blending to the outer muscle look. A combination of Maya plug-ins was used to generate and animate the muscle fibers. The fibers were then converted to RenderMan curve primitives. Using a RenderMan DSO, we could determine how far below the muscle's surface each fiber lay and start blending to the surface based on this distance. This gave the effect of the fibers gradually condensing into a solid muscle.

The transformation sequence pushed most of our preconceived ideas about rendering to the limits and forced us to think up new strategies to overcome the problems that we were presented with. We not only had the challenge of making CG models look like living flesh, but we also had to extend this look into the vacuum inside our CG models. Using a multitude of ideas, techniques, and tools, we were able to do this and also render it in a feasible amount of time.

For “The Legend of Sleepy Hollow,” Tim Burton created a dream sequence where the lead character’s mother drops out of a medieval Iron Maiden toward the camera. It’s a very shocking and disturbing sequence to watch. The shot worked well as it was originally shot, but Mr. Burton wanted more dramatic impact and asked Industrial Light + Magic to take on that challenge. Could we add blood pouring out of an Iron Maiden, rushing toward the camera, as the door of the coffin-like Iron Maiden opens?

Industrial Light + Magic had originally planned to shoot all the blood elements practically, then composite them into the shot digitally. Jim Mitchell, the effects supervisor of the film, needed some kind of automatic to fully formulate the concept.

My first job on the team was to investigate off-the-shelf solutions that would help generate ideas for the composition. The more we worked on the project, the more Jim Mitchell became convinced that it would be just as difficult to create digital blood as it would be to create the elements practically. It was very difficult to develop convincing blood. Particle-based systems would allow us to control the animation exactly as we wanted for compositional reasons; however, the dynamics just didn’t behave like blood.

John Anderson, who developed various physical simulation engines for “The Phantom Menace,” and Chris Horvath, who first experimented with simulating splashing liquids on “Deep Impact,” were developing liquid simulation software to be used by several ILM projects. In addition to this, Florian Kainz was building a new volume renderer with Chris assisting in the test and development phase. All together, we were able to generate some very beautiful looking volumetrically rendered blood.

When Jim Mitchell saw John and Chris’s blood simulation tests, he was confident we could use their software for this project. Nevertheless, we still had time to create practical elements should we get into trouble with the new software. John and Chris worked hard for several weeks developing the software and integrating it into our pipeline.

Once we got the software, the next step was how to control it and how to control an artist, such as myself, with this new tool. Maya became the main interface for developing the animated geometry environment; the blood simulation was developed with ILM’s own proprietary system, called Hybrid.

Chris and I worked very closely, testing newer and newer software versions until Chris and John were able to stabilize the product. Once it was stabilized, they continued to upgrade the software right up until the shot deadline.

Tim Burton loved it. It was great seeing it in the movie theater and listening to the audience’s reaction. Mission accomplished.

This sketch explains some of the technical aspects of our blood simulation software and new renderer. Through a variety of video clips, we demonstrate the procedural animation development process and the well-documented steps that led to the final result.

This sketch describes a process for automatically matching 3D models to sequences of hand-drawn animation. This work was undertaken with the idea of ultimately modeling, tracking, and conforming 3D sets, characters, and props to traditionally animated drawings. Here we outline the techniques developed so far to achieve this goal with simple rigid props.

We have constructed a plug-in for Maya, a commercial modeling/animation package, to assist in the process of integrating 3D CG elements into traditionally animated films. Our plug-in was inspired by computer vision work in which CG rendering is used in a feedback loop (“visual servo”) to iteratively optimize parameters of camera or object models.¹

Our process begins by coloring the cel animation with a special “tracking palette” of clearly distinguished colors. We render our 3D model of the object in the same special set of colors. We initialize an iterative numerical optimization by constructing a scene and camera to roughly position the object in alignment. The rendered scene is compared to the specially colored animation cel (the target image), with the error metric calculated as the sum of the squared pixel differences. The optimizer drives the object’s pose by iteratively perturbing transformation parameters, re-rendering the scene, and recomputing the error. Progression stops when the error cannot be further reduced. This process is completed for each frame, progressing through the sequence of images for a scene. Using the solution at each frame as a starting point for the next increases the efficiency of the process and the continuity of the solution.

Typically, no static shape and set of camera parameters can accurately describe a hand-drawn object for more than a few frames of animation. Yet it is imperative for our work that we register the prop silhouette exactly to the lines drawn by the animator. To that end, we introduce an automatic 2D warping step. Computing an implicitization of the silhouette of the target image yields a smooth, signed pseudometric of distance from the original outline. Overlaying the silhouette of the best rigid body fit, we are able to deterministically walk points on the silhouette of the 3D object to the pseudometrically closest points on the drawing’s outline, establishing a set of vectors that smoothly warp the outer edge of the object to achieve perfect alignment with the drawn target image. Membrane interpolation distributes this warp evenly over the object interior.

We show two results. The first is tracking a drawn sequence of a prop and replacing it with a shaded 3D model. Even in the case of this simple bowl, the animation was “cheated” for dramatic effect, and our 3D model bowl had to scale smoothly in the vertical dimension to match the drawings. The second sequence, of a hand-held torch, took more dramatic license, and it was necessary to deform the object with the 2D warping step in order to match the animation. An additional sequence shows the torch replaced with a more complex prop, which illustrates the editorial latitude supported by our technique!

Reference

1. Busch, H. (1989). Automatic modeling of rigid 3-D objects using an analysis by synthesis system. *Proc. of the SPIE Conference on Visual Communication and Image Processing IV*, 1199, 356ff.



The Pharaoh’s magician’s Hotep and Huy from DreamWorks Pictures’ animated feature “The Prince of Egypt” hold a 3D bowl during the magical sequence “Playing With the Big Boys.”



The Pharaoh’s son holds up a replaced 3D torch in DreamWorks Pictures’ animated feature “The Prince of Egypt.”

There are two main problems in the metamorphosis of two 2D shapes: vertex correspondence and vertex path problems. The solutions to the vertex path problem are presented by Carmel & Cohen-Or¹ and Sederberg et al.² The shortcoming of these approaches is that they only consider the interpolation between two object's boundaries. Here, we present a solution to the vertex path problem when the source and the target shapes contain more than one piecewise linear curve. For simplicity, we use the term "curve" in place of piecewise linear curve.

PROPOSED VERTEX PATH ALGORITHM

Let $S = \{C_1^s, \dots, C_n^s\}$ and $T = \{C_1^t, \dots, C_n^t\}$ to be the source and the target shapes. Here, C_i^s (C_i^s or C_i^t), $i = 1, \dots, n$ is a curve. Let m_i to be the number of vertices in C_i^s . Then $C_i^s = \{v_{i,1}^s, \dots, v_{i,m_i}^s\}$, where $v_{i,j}^s$ is a vertex. We assume that $v_{i,j}^s$ corresponds to $v_{i,j}^t$. Our algorithm consists of a preprocessing step and an intermediate shape generation step.

PREPROCESSING STEP

In the preprocessing step, we construct a dependency graph. The dependency graph is constructed by generating a connection graph and then traversing the connection graph starting from the base vertex in a breadth-first manner and obtaining the intrinsic definitions² of the edges.

We can consider curve c_i^s as an undirected graph. We call this the "basic connection graph." The nodes and edges of the basic connection graph represent the vertices and the line segments in the original curve. We build the connection graph by linking the basic connection graphs that are located close to each other by putting edges among them.

Figure 1(b) shows the dependency graph with respect to the connection graph in Figure 1(a). To build the dependency graph, we first create a root node for the base vertex. A vertex that minimizes $\|v_{i,j}^s - v_{i,j}^t\|$, $i = 1, \dots, n, j = 1, \dots, m_i$, is chosen as the base vertex. In Figure 1(a), we assume that $v_{2,2}$ is the base vertex. Next, one of the edges that is attached to the base vertex is chosen as a reference edge. For the connection graph in Figure 1(a), $v_{2,2}v_{2,1}$ is chosen as the reference edge.

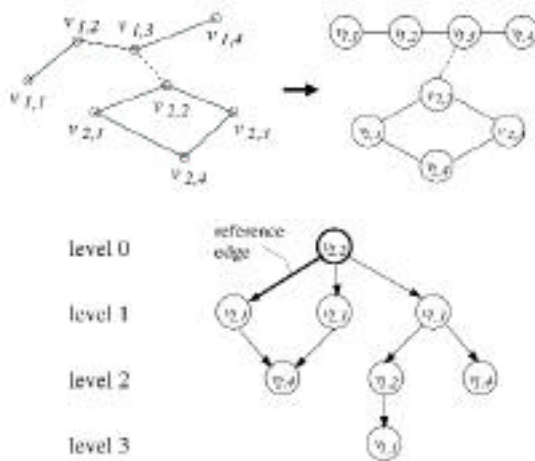


Figure 1: Examples of connection graph (a) and dependency graph (b).

We calculate the intrinsic definitions of the reference edge with respect to the x axis. The intrinsic definitions of the remaining edges, which have the base vertex as one of their end points ($\overrightarrow{v_{2,2}v_{2,3}}, \overrightarrow{v_{2,2}v_{1,3}}$), are calculated with respect to the reference edge. We then traverse the connection graph starting from the nodes ($v_{2,1}, v_{2,3}, v_{1,3}$) that share edges with the base vertex in a breadth-first manner and assign a direction to each edge. When there is an edge from $v_{i,j}$ toward $v_{k,l}$, we store the intrinsic definitions of edge $\overrightarrow{v_{i,j}v_{k,l}}$ of the source and the target shapes in node $v_{k,l}$. We use these intrinsic definitions in the edge-angle blending computation. For instance, in Figure 1(b), node $v_{1,2}$ holds the intrinsic definitions of edge $\overrightarrow{v_{1,2}v_{1,3}}$ with respect to edge $\overrightarrow{v_{2,2}v_{1,3}}$. Note that a node can store more than one intrinsic definition and thus can have more than one dependency.

GENERATION OF INTERMEDIATE SHAPES

The intermediate shape at time t is computed as follows. Let B^s and B^t be the coordinates of the base vertex in the source and in the target. We compute the coordinates of the base vertex B^l by using the linear interpolation of B^s and B^t . Then, we traverse the dependency graph, starting from the root in a breadth-first manner, and compute the coordinates of the vertices in the order visited. Suppose that vertex $v_{-(i,j)}$ has d dependencies. We compute the new position of $v_{i,j}$ by using $v_{i,j}^l = \sum_{k=1}^d v_{i,j,k}^l / d$. Here, $v_{i,j,k}^l$ ($1 \leq k \leq d$) are the new positions of $v_{i,j}$ with respect to each dependency. We calculate $v_{i,j,k}^l$ using the edge-angle blending method.²

EXAMPLES

Figure 2(a) shows a result of our vertex path algorithm. There are some parts in the source and target shapes that have different orientations. Nevertheless, our algorithm produces a smooth interpolation sequence. The relative positions of the eyes, the nose, and the mouth are preserved during the interpolation process. We implemented this vertex path algorithm as a part of our morphing system. Figure 2(b) shows an example of morphing from a bear to a gorilla.

References

- Carmel, E. and Cohen-Or, D. (1997). Warp-guided object-space morphing. *The Visual Computer* 13, 465-478.
- Sederberg, T., Gao, P., Wang, G., & Mu, H. (1993). 2D shape blending: An intrinsic solution to the vertex path problem. *Computer Graphics (Proc. SIGGRAPH 93)*, 15-18.




Figure 2: Results of our vertex path algorithm (a) and its application to morphing (b).

The objective of our research is to investigate whether it is possible to construct a fast and inexpensive parallel rendering system leveraging the aggregate performance of multiple commodity graphics accelerators in PCs connected by a system area network.

We investigate a sort-first approach.^{1,2,3} As shown in Figure 1, our system consists of a client PC, server PCs, and a display PC arranged in a three-stage pipeline. During the first stage, the client partitions the screen into P tiles and assigns each of them to a different server. During the next stage, each server renders all the graphics primitives at least partially overlapping its assigned tile from a replicated scene database to form a subimage and sends the resulting pixels to the display PC. Finally, the display PC composites all the subimages in its frame buffer for display.

The main challenge is to develop a dynamic screen-partitioning algorithm that balances the rendering load among the servers, minimizes overheads by reducing primitive-tile overlaps, and executes at interactive rates.

Our method is based on the mesh-based adaptive decomposition (MAHD) algorithm described by Mueller.² Since it is not possible for the client PC to transform and sort every graphics primitive among the tiles, we group primitives into coarse-grained objects. The client partitions the screen based on the predicted rendering costs of these objects, while the servers perform detailed sorting for each image using the bounding volume hierarchy as they render objects. The key idea is that the client and servers share the costs of sorting graphics primitives among tiles, with the major portion of the sorting done in parallel by the servers.




Figure 1: System architecture.

To test the effectiveness of our sort-first approach, we have implemented a prototype system using a cluster of 18 Pentium III PCs equipped with \$250 Hercules 3D Prophet graphics accelerator cards. The PCs communicate over a Myrinet network, which delivers approximately 100MB/s peak bandwidth with 13μs one-way latency. Results from our experiments show that it is possible to achieve good speedups (8.1) on a cluster with up to 16 server PCs (see Figure 2). Surprisingly, the limited communication performance of the PC cluster is not the major bottleneck, but rather traditional overheads of parallel rendering limit the scalability of our system. In particular, client processing time grows linearly, server overheads grow linearly due to increasing primitive-tile overlaps, display PC bandwidth becomes a bottleneck beyond 25 frames per second, and the memory capacity of each server limits the size of the 3D model. Alleviating these limitations is a focus for future research.

Overall, we conclude that PC clusters provide an attractive, lower-price alternative for high-performance rendering. By building a working prototype system, we have demonstrated that a sort-first algorithm can work for moderate numbers of processors, and that a PC cluster architecture can be an effective way to combine multiple graphics accelerators together.

References

1. Molnar, S., Cox, M., Ellsworth, D., & Fuchs, H. (1994). A sorting classification of parallel rendering. *IEEE Computer Graphics and Applications*, 14 (4), 23-32.
2. Mueller, C. (1995). The sort-first rendering architecture for high-performance graphics. 1995 Symposium on Interactive 3-D Graphics.
3. Samanta, R., Zheng, J., Funkhouser, T., Li, K., & Singh, J.P. (1999, August). Load balancing for multi-projector rendering systems. SIGGRAPH/Eurographics Workshop on Graphics Hardware.




Figure 2: Timing results gathered during tests with a polygonal model representing a skull from the Visible Human Project (2,111 objects and 355,511 polygons). The height of each bar represents the time (in milliseconds) taken by each of our three pipeline stages. Server bars are split into two parts: productive rendering (below) and overheads (above). Speedups are indicated above each bar.

The Mixed Reality (MR) stage is a networked layering of physical space and virtual environments. It enables multiple performers to interact and communicate through intuitive free-body, non-contact interfaces. Connecting the concept of the stage with the idea of digital information space comprises investigation in digital storytelling and design of non-linear structures. This sketch reviews the concept of a MR stage, technological implementation, and strategy for development; experiences with public trials; lessons learned “from the real world;” and questions for further development.


The metaphor we use for the MR stage is that of a room furnished or filled with data. The “room” represents the physical interaction space, but the furniture of data is virtual and represents a 3D information space. It is spatially organized information in which audio-visual data are revealed through the user’s movement in the combined real-virtual (MR) space and through interaction with other users.

The technical system developed for the concept of the MR stage is eMUSE, an electronic Multi User Stage Environment, based on VRML 97. It is built as a modular system that provides independent levels of implementation for interfacing, rendering, and displaying the virtual environment; for supporting multiple users in shared physical and virtual space; and for supporting attachment of non-standard interfaces (for example, a vision system). By separating the browser and the interface, eMUSE enables attachment of arbitrary interfaces to the normally closed VRML browser.

The MR stage works through interplay of the following elements:

- The vision system interface.
- The virtual environment (VRML) audio-visual content with the user as avatar-trace
- The physical interaction space.

Unlike in systems where navigation in virtual space is achieved by direct manipulation of some physical device, the MR stage cannot be navigated without the existence of a congruent physical space. The real body becomes an essential part of the MR interface environment.



We investigated the usability of the system as an interactive installation at the Berlin media art festival Transmediale'99, with 750 accidental visitors. During the Bochum Fiden'a'99 theatre festival, we explored the conditions of the MR stage in theatrical space with two professional performers.



Experiences from these public trials point to the following scheme for the perceptive process in an interactive environment:

1. The user tries to identify the structures and rules.
2. The user plays with the structures and rules
3. The user considers what actions look like and how they sound to the audience.
4. The user becomes aware of the other player(s).
5. The user tries to communicate with the other(s).

The Mixed Reality installation created the impression of a cinematic space controlled by the visitors. In the theatrical situation of the MR stage, the performer appeared as an instrument touched by invisible sounds.

“Superbodies”

Alexander Tsiaras
Anatomical Travelogue, Inc.
atsiaras@anatomicaltravel.com

National Geographic Television
Pangolin Pictures, Inc.

The Anatomical Travelogue has signed the largest contract in National Geographic Television’s history to produce a series of programs over several years that map the human body, from the smallest cells to the gross anatomy.

Using real human data, the series will illustrate human performance at its extreme. In the incredible “Superbodies” journey, we examine the human organism at the peak of performance. We digitally scan the world’s greatest athletes, artists, and scientists to show quantitatively accurate human data in ways no TV camera ever could.

To site only a few examples, the images will record; the world’s greatest free diver’s respiratory system dynamically recorded during a record dive, a woman’s endocrine system activated to increase endurance of pain during delivery of a baby, and a great comedian’s quick-witted brain recorded in a functional MRI.

Working closely with Volume Graphics to produce proprietary software to complete the series, Anatomical Travelogue will be acquiring a large group of the most important datasets of quantifiably important anatomical, physiological, biological, and bio-mechanical data ever investigated.



Survivors: Testimonies of the Holocaust: Humanity through Technology

Stephanie Barish
University of Southern California
sbarish@annenberg.edu

Scott Chamberlin
Shoah Foundation

“Survivors: Testimonies of the Holocaust,” a recent CD-ROM from the Survivors of the Shoah Visual History Foundation, merged the practical opportunities of multimedia technology with intense humanistic content. As a result, it serves as a model for what might truly be the successor to the conventional textbook, using interactive technology not just as a means of delivery, but also as a subject-appropriate device that brings students and history together in original ways. The CD-ROM has advanced the craft of interactive presentation by breaking down perceived “barriers” between technology and humanity.

Between the beginning of R&D in 1995 and the CD-ROM’s completion in January 1999, mainstream interactive technology developed from a medium in its infancy to one that could bring sustained video, full-screen animation, and nearly unlimited hypertext interactivity to common household (and “schoolhouse-hold”) computers. Meanwhile, the developers worked with focus groups to try to understand the damaging perception that digital multimedia is incompatible with a sense of humanism, discussing such factors as the perceived “illegibility” of on-screen text, the forbidding “digital” appearance of interfaces, and the common perception that interactive technology was ultimately more frustrating than it was exciting.

The overall goal was to create an appropriate and respectful interface for complex Holocaust-related materials that included video footage of survivor testimony. Through this process, a technology-friendly


audience of students could be reached, and a largely technophobic audience of schoolteachers could be “converted.” In addition, the team wanted to use the interest generated by survivor testimony to encourage reading of significant written content — essentially, to use natural curiosity to “trick” students into reading more written materials than they might otherwise care to do.

Technological advancement and an extended production allowed the developers to address their specific challenges through unique means. Primary among them is the painstaking rendering of thousands of pages of text: every caption, description, and page of historical text was created graphically to improve the readability and appeal of the words. The main interfaces, too, deliberately suit the video, photographic, and textual elements of the overall presentation. Not least, the interactive structure carefully emphasizes the chronological nature of the materials, encouraging curious deviation from the survivors’ stories while ensuring that the user follows the materials in an educationally valuable way.

This sketch presents a discussion on the use of interactive technology as an educational tool that can truly bring people closer to information and to one another.



Four interfaces from “Survivors: Testimonies of the Holocaust.” Clockwise from top left: Archival images illustrate a survivor’s memory of her hometown; February 1941 Holocaust Timeline image shows graphically rendered historical text with hypertext links and roll-over caption; a “map detail” screen shows a survivor’s account in a geographical interface; the “reference library” page on Joseph Stalin with path buttons.



Front cover for “Survivors: Testimonies of the Holocaust.”

Cultural heritage can benefit greatly from the use of virtual reality and its integration with hypermedia. The advantages of using these technologies include reconstruction of buildings, sites, or objects that no longer exist or that exist partially; simulation of viewpoints that are not available in the real world; and direct interaction with virtual reproductions of works of art and the consequent elimination of the risk of degradation. Virtual reconstruction of objects of cultural importance that in the real world are not in the same place, or the evidence of which is only fragments and/or historical documentation, offers insight and stimulates interest and curiosity.

There is no doubt that different users (tourists, students, researchers, etc.) have different interests and needs: for example a researcher needs a more detailed visualization than a tourist, while a tourist could be interested in general information about buildings that are not interesting to historians because they are well-known. The aim of this project is to develop an innovative interface system (implemented with a specific XML-based markup language) for accessing a corpus of cultural heritage data that are coded in order to allow the user to personalize the type of information to be received during a virtual tour. Users access a tailored version of the virtual environment, which is built dynamically and interactively on the basis of their language, profile, specific interests, level of competence, and knowledge. The basic idea is to enable the user to choose, through the interface, which elements are to be visualized during the virtual tour. In this way, users can visualize only the virtual components that interest them (for example, architectural elements, decorations of a particular period or style,

recontextualization of objects — sculptures, tools, frescoes, floor tiles, mosaics, etc. — that have been removed for exhibition in museums or have been transferred to depositories. In order to provide the visitor with “interesting” historical or archaeological information, every 3D object is linked to a hypertext description of the resources on which the virtual reconstruction is based. Also, users can personalize the type of information they want to receive during their tour. For example, they can access tourist information, simple descriptions of the 3D objects, or more details about related resources or the archaeological hypothesis.

We are developing a prototype Web VR application based on a case study of multimedia data selected to explain some archaeological finds in Pompeii, the most complete Roman town in the world. The study, a joint project of CINECA and the Bologna University Department of Archaeology, focuses on two topics: the insula IX, of Pompeii and a corpus of mythological images in domestic spaces (Hercules Domesticus).

An high-resolution version all the reconstructed 3D objects can be visualized in a virtual environment enjoyable at the CINECA Virtual Theatre (www.cineca.it/visit/virtualtheatre.html).

The authors wish to thank Daniela Scagliarini and Antonella Coralini, Archaeology Department, Bologna University; Tullio Salmon, DEIS Bologna University; the Soprintendenza Archeologica di Pompeii; all the members of the HYLAS Project; Claudio Alvisi, CINECA; and Maria Teresa Pareschi, CNR Pisa National Research Council.



Study for a virtual environment to access archaeological multimedia data.

“Tal Vez Mañana” is the result of the research and production of Medio y Medio Productions and explores the capacities of multimedia language. But it is also more than that.

The audiovisual interface for the narration of the four stories in “Tal Vez Mañana” is based on a strategy that enables the hypermedia reader to quickly and intuitively recognize and adopt new reading codes, thus assuming an active role in the text. The reader can go through each story sequentially or interact as the interface indicates, thus breaking the linearity of the text. It is possible to navigate among narrative elements, choose the next link of the story, and determine a narrative path. In this way, you can arrive at a particular scene within the same narrative but seen from another point of view. The narrative threads are distinguished by aesthetic and audio-visual design, as well as by the usage of particular language within each of the four stories.

TVM is an interactive novel, a product of our own culture, imbued with the popular and the autochthonous. The narratives of this work are enabled by new media and the new technologies of globalization, yet the content is particular and local. This work exhibits a novel use of hypermedia and a search for new ways of telling stories.

A real environment in Montevideo, the city where the TVM story develops, is filled with everyday events, anger, love, and disappointments. A boy (Carlitos), a taxi-driver (Jesus), a lame woman (Nahir), and a waiter (Ramsn) are the main characters. Among the tango and the candombe, Ramsn expresses his misfortune to the patrons in the bar. Jesus and Nahir fall in love in the Temple of Venus, while Carlitos attempts to retrieve a ball that dropped into the fountain of the Temple. They get into the taxi-driver’s car as he starts his daily routine, but he is unaware of his coming new love, and deception.

The design of each story is in a different style.

Carlitos: This segment uses drawings collaged onto photographs, using color to represent the boy’s emotions and the atmosphere of the environment. Movement is shown through simple animation, reminiscent of Japanese anime and Manga design.



Nahir: The photo-novel is used as code structure for the construction of this segment; multimedia elements are used for background sounds and dialogue. Screen images are built from different shots within an environment, and of the characters through the use of layers.

Taxi-Driver: Paintings in watercolor were digitally retouched and animated in Director.

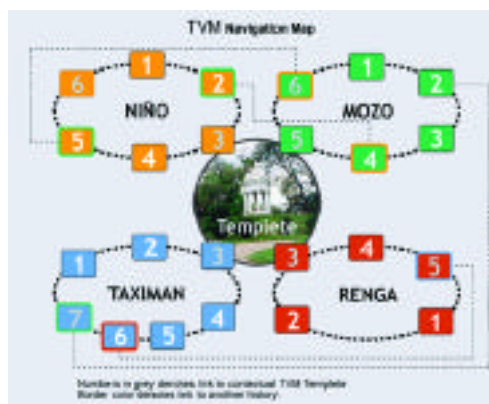
Ramsn: A sequence of three images, with a primary central image. This sequence is based on animation of some frames extracted from video or digital photography and composed in actual environments, while others were constructed using a blue screen.

Based on the reader’s navigation, the narration goes on through the four stories sequentially, or it can jump from one story to another. The central environment of the Temple of Venus is the navigation engine, and the starting point of the novel, where linearity enables the reader to use clear codes and become aware of the level of interaction as the narrative is engaged.

For construction of the audio-visual interface space for the dialogue between the reader and the novel, we referred to the historical audio-visual languages for which codes are already culturally absorbed. The novel’s strategy, its stages, characters, and elements allow the reader to experiment with nonlinear navigation. The interface design also results in a visual syntax that derives from the elements of graphic design, which come together on the screen.

Where do we find the story of TVM? In each narrative thread, within a group of them, or in all of them. The experience of each reading allows this literary work to construct a meta story with new capacities for reading, narrative construction, and involvement of the reader in the story. The authors also faced challenges: creating texts, weaving complex plots, creating a space where a new construction of meaning is possible in each reading, and designing a special freedom for the reader.

Polyphony resounds in “Tal Vez Mañana,” in its characters, environments, multiple readings, in its potential constructions, and in its complex and non-hierarchical structure.



Texture synthesis is an interesting and important problem in the field of computer graphics. Recently, several techniques have emerged in the literature that are able to analyze an input texture sample and synthesize many new random similar-looking textures.^{2,4,5} This work extends these techniques in several important ways. First, we describe a new texture-mixing algorithm, a statistical learning algorithm that operates on several different input texture samples to synthesize a new texture. This new texture is statistically similar to all of the input samples, and it exhibits a mixture of their features. Second, we extend our approach from the domain of static textures to the domain of texture movies: dynamic, time-varying textures, or TVTs for short. More specifically, we present a technique capable of generating a sequence of frames, corresponding to temporal evolution of a natural texture or pattern, that appears similar to an input frame sequence. For example, using this technique, we generated short movies of dynamic phenomena, such as waterfalls, fire flames, a school of jellyfish, turbulent clouds, an erupting volcano, and a crowd of people. The generated sequences are distinguishable from the original input sample, yet they manage to capture the essential perceptual characteristics and the global temporal behavior observed in the input sequence. A specialized variant of our technique is able to synthesize one-dimensional sound textures, such as sounds of traffic, water, and certain kinds of music.¹


The natural applications of texture movie synthesis are in special effects for motion pictures and television, computer-generated animation, computer games, and computer art. Our method allows its user to produce many different movie clips from the same input example. Thus, a special effects technical director could fill an entire stadium with ecstatic fans from a movie of a small group, or populate

an underwater shot with schools of fish. Designers of 3D virtual worlds will be able to insert animations of clouds, smoke, and water from a small number of input samples, without ever repeating the same animation in different places. However, these are by no means the only applications of such a technique. For example, methods for statistical learning of 2D texture images have been successfully applied not only to texture synthesis, but also to texture recognition and image denoising.³ These applications are made possible by realizing that statistical learning of 2D textures implicitly constructs a statistical model that describes images of a particular class. Similarly, our approach for learning TVTs can be used as a statistical model suitable for describing TVTs. Therefore, it should be possible to apply this statistical model for tasks such as classification and recognition of such movie segments.


Additional information can be found on the SIGGRAPH 2000 Conference Abstracts and Applications CD-ROM and at:
www.cs.huji.ac.il/~danix/research/texsyn/index.html

References

1. Bar-Joseph, Z., Lischinski, D., Werman, M., Dubnov, M., & El-Yaniv, R. (1999). Granular synthesis of sound textures using statistical learning. *ICMC 99 Conference Proceedings*.
2. De Bonet, J. S. (1997). Multiresolution sampling procedure for analysis and synthesis of texture images. *ACM SIGGRAPH 97 Conference Proceedings*, 361-368.
3. De Bonet, J.S., & Viola, P. (1997). A non-parametric multi-scale statistical model for natural images. *Advances in Neural Information Processing*, 10.
4. Heeger, D. J. & Bergen, J. R. (1995). Pyramid-based texture analysis/synthesis. *ACM SIGGRAPH 95 Conference Proceedings*, 229-238.
5. Zhu, S. C., Wu, Y., & Mumford, D. (1998). Filters, random fields and maximum entropy (frame) - towards a unified theory for texture modeling. *International Journal of Computer Vision*, 27 (2), 107-126.



Texture mixing: The larger texture on the right was synthesized by mixing the two smaller textures on the left.



Texture movie synthesis: The three images in the top row are frames from an input waterfall sequence. The three images in the bottom row are frames from a synthetic waterfall sequence.

Art in public spaces has almost exclusively been a static phenomenon, where a single artwork is displayed over a very long period of time. The progress in new media and information technology permits new dynamic approaches to public art, where the content may evolve continuously, and interactivity and connectivity may play an important role.

We have built a system that uses the latest advances in projection techniques together with a custom-built computer system to create an experimental arena in a public city-space, where artists working in new media can take part in the on-going creation of a ever-changing content. The ambition of the project has been to create a new meeting place in the city, and to provide it with a dynamic and interactive artistic installation. The component of the project described in this sketch displays computer-generated images simultaneously on four sides of a glass cube of human proportions.

The project, called Time Document, is the official Swedish installation/monument intended to commemorate the millennium. The project was actually initiated by the Swedish king, Carl XVI Gustaf, who in contacts with the millennium committee of the Swedish government expressed a wish for a more permanent millennium installation. The entire project was carried out in a collaboration among architects, artists, and scientists at Chalmers University of Technology. The installation is intended to be permanent and to be continuously supplied with new content. It has been in operation since the beginning of 2000.

The installation is located in central Göteborg, Sweden, in an open space next to the old opera house (Stora Teatern). Göteborg is the second-largest city in Sweden, with a population of about 600,000. A very large number of people pass the installation on foot each day, especially on weekends. The space itself can be viewed as an abstract garden, where a visitor can walk among several interesting objects, or sit down and rest. In addition to the cube, this open space contains a number of other objects, in particular a tall sculpture containing a computer-controlled light installation (where light is transmitted via optical fiber from 40 individually controlled light sources) and a sound sculpture based on real-time-generated, never-repeating algorithmic compositions.

The artistic contributions in the cube fall into four categories:

1. Pre-recorded four-channel video material, stored as MPEG-2.
2. Algorithmic computer art and real-time simulations.
3. Interactive artistic experiments that use image processing for interactions with viewers around the cube.
4. Experiments involving interactivity over the Internet.

In addition to artistic contributions, the cube shows material based on visualizations of scientific phenomena. Several artistic experiments involving interactions over the Internet are currently in progress. As an example, Henrik Hokansson, an internationally recognized Swedish artist known for work that involves biological systems, is developing an installation connecting the cube to an underwater site in the North Sea outside Göteborg. A number of projects also involve algorithmically generated images, and an international competition for algorithmic art is also planned in connection with the cube.

TECHNICAL DETAILS

The part of the system that is visible above ground consists of a glass cube with a side of two meters, which is raised about 50 cm from the ground (See Figure 2). The size of the cube was chosen so that it can display human figures in natural size. The glass cube consists of several laminated layers, with a single layer of a film suitable for back-projection sandwiched in between. The risk of vandalism was a major concern in the design of the cube, but so far there have been no incidents.

The cube has five glass sides (a ceiling but no floor). Images are projected from the inside on four sides. Two speakers are located in each side of the fundament. Below ground, four powerful video projectors (with 1280 x 1024 pixels resolution) and part of the computer system controlling them are located in a climate-controlled 36 m² room. A section of the room and the cube on top of it are shown in Figure 2 to illustrate the geometry of the system. Most of the perspective correction is done optically, but this has to be supplemented by a smaller software correction because of the extreme projection angle. Several image formats are available, including a mapping from 1280 x 1024 pixels to a square image, and a mapping restricting the image to 1024 x 1024 pixels.

The cube is driven by a Linux-based system that currently consists of six networked double-processor PCs. The system design and all software development were carried out by ourselves. The choice of a Linux system was motivated by reliability and safety issues, as well as economic reasons. Videos are stored as MPEG-2 on an array of hard drives for each channel and decompressed in hardware. Programs generating graphics in real time can run in parallel on four machines, one per channel, with synchronization and communication of boundary conditions over the network. The system is equipped with one camera per side, for interactivity through image processing. The system is equipped with a fast network connection to the university network, which in future experiments, for example, will be used to connect the system to a VR cube at Chalmers.

This project was supported by the the Knut and Alice Wallenberg Foundation, the Foundation Culture of the Future, and the Erna and Victor Hasselblad Foundation.

ABSTRACT

This sketch introduces the traditional animation concept of the timing chart and illustrates how these abstractions can be used to simplify creation and spacing of keyframes in an animation package.

BACKGROUND

Typically, a computer animator needs to sift through and manipulate tens, hundreds, and sometimes thousands of animation curves in order to produce a typical piece of character animation. The animator's traditional tools to manipulate this data, a graph editor or dope sheet, can present the user with an overwhelming amount of information. These tools display data in a non-intuitive fashion, making it difficult for our traditional directors to give timing notes to animators and making the job of transitioning traditional animators into computer animators problematic. Timing charts, by contrast, serve as a job management tool, allowing an animator to initially focus on performance, postponing timing decisions until later in the process.

TimingChart is an extension to the ImageTimer animation system first presented in the SIGGRAPH 98 animation sketch "ImageTimer-A Traditional Approach to 3D Character Animation."

OVERVIEW

Our approach to timing animation is taken from the traditional animation world — an approach that is used to this day by every animation house to time every scene of animation. Once a scene's key poses are drawn and a timed pose test is created, the animator creates hand-drawn timing charts that instruct an assistant animator how in-betweens should be drawn between the key poses. While this step seems unnecessary in the digital world where the computer performs in-betweening, timing charts provide an efficient control mechanism.

So what is a timing chart and how does it work?

Assume we have two frames of a falling ball animation. The first drawing occurs at frame 1, the second at frame 7. Diagram 1A indicates what the timing of the ball would be with no in-betweens ("tweens"). Diagram 1B indicates that a tween halfway between frames 1 and 7 (indicated as the blue ball) should fall at frame 5. A timing chart can define an arbitrary number of tweens between two key poses. Note that these simple diagrams give an animator an exact specification of both the position and orientation of the pose and its timing.




Figure 1: How in-betweens work.

GETTING STARTED

Let's see how we can translate the ideas of hand-drawn timing charts into the computer-animation workflow. First, we create a chart by binding a set of animation controls to a chart name. Once this association has been made, the animator can begin creating poses and timing. TimingChart supports two modes of operation: Key/Tween mode and Arc mode. The first mode is used in setting and deleting keys and tweens. The second mode is used in creating and deleting timing arcs defined between keys. Timing arcs provide a graphical representation of the timing between keys. Once defined, tweens are placed at arc endpoints, and their corresponding frame positions are defined. The Body column of Figure 2 is an example of a simple timing chart.

By default, a tween's associated frame number is based on the pose ratio defined by neighboring keys and tweens. For instance, frame 6 is 1/4th of the way between frame 5 and 9. The frame number of a tween can be adjusted easily. Note that not every arc end point need have an in-between defined.

CREATING A SUB-CHART

While a chart for an object might provide enough refinement to animate a rigid body or simple background character, high-quality animation demands finer grain control to impart anticipation, overlapping motion, and follow-through to characters. In the traditional world, animators construct partial drawings ("partials") to help achieve subtle timing differences within a character. TimingChart has a parallel notion of partials called branches that can be created from within a sub-chart.




Figure 2: Sample and sub-chart.

A sub-chart is defined as a subset of the objects associated with any given chart. Sub-charts can not share objects between them. Sub-charts allow an animator to override timing over select frame ranges, through creation of branches. The Left Arm column of Figure 2 illustrates a sub-chart and single branch created to override the timing of the left arm from frames 5 through 13.

EXAMPLE

TimingChart was used to create the after-hours film "Oil & Vinegar." Every scene was animated with TimingChart, and the graph editor was rarely needed. A sample image is shown below.



Figure 3: Image from "Oil & Vinegar," created with TimingChart.

In the making of “Toy Story 2,” we had several technical challenges, mostly dependent on the amount of geometry and complexity of every model and associated shader. This sketch addresses three areas where we had to develop custom solutions in a very short amount of time:

- *Hair for the humans and for the dog Buster.* We decided to design a new set of tools for hair sculpting and generation to allow us to produce hair and fur in a short amount of time. We wrote several plug-ins for commercial applications to allow our modelers to sculpt hairs and custom code for RenderMan for the generation process. The modeling process is completely controlled inside a 3D modeling program. After defining the regions where the hair would grow, one or more subdivision surfaces were created. These surfaces represent the scalp. From these layers, the modeler can sculpt a set of “key” hair and let the procedural model compute the “interpolated” hairs. A custom tool was developed to control various noise and deformation parameters used to achieve the final look.
- *Procedural dust generation.* To address the visual complexity of a “dusty” environment, we decided to use a procedural system that generates actual geometry on any given surface. Early experiments with surface shaders demonstrated the ineffectiveness of this approach, especially for extreme closeups on the “dusty” surface. Using a custom plug-in that converts any modeled geometry into

C code, we wrote a RenderMan DSO module that used texture maps for controlling distribution and height, and was able to scatter several dust particles on multiple surfaces. A camera was placed on top of all the surfaces with dust. A first RGB and Z render pass was done from that camera. Using a paint program and gray-scale maps, we were able to control the density of the dust. The procedural module then used these maps to place dust geometries.

- *Use of Level of Detail.* For the final sequence in the airport, we had to use several optimization techniques, since we were getting close to the physical limit of our renderer: two GB of RAM allocated for a single frame. Also, rendering time for a single frame was getting prohibitive for film-resolution renders. To overcome these limitations, we enabled the Level of Detail and delayed the loading feature implemented in RenderMan. For every surface, we provided its bounding box and let the renderer compute the area in pixels of that bounding box on screen. If it was smaller than a pixel, we faded the surface away. This allowed us to remove an incredible number of details on geometries rendered in background. We also bounded complex models and told the renderer to load them only when required and when they were visible on screen.

This sketch introduces 3D use of 2D drawn-water effects animation as seen through a 3D camera while being played back in 3D space on digital cards driven by 3D particle systems to achieve complex environmental effects for “The Road to El Dorado.”


Doug Ikeler, sequence lead on the film, advanced a theory first tried on DreamWorks SKG’s first animated film, “The Prince of Egypt.” He wanted to try using traditionally drawn artwork effects for animation cycles attached to, triggered, and driven by 3D particle animation in ways that would allow creation and production of many scenes involving crashing whitewater. He created a tool set he called Spryticle to do the complex, rapid water effects in the climactic sequence of “The Road to El Dorado.”

The particle animation uses flat digital cards that carry triggered playback of scanned traditional 2D animation. Doug received hand-drawn cycles of black-and-white water splashes that had been visualized from specific angles. After they had been cleaned up, scanned, and digitally regioned in the DreamWorks Animo system, Doug brought them into Maya as images to be displayed on cards in 3D space. The flat digital cards in Spryticle can be bent, curved, stretched, and positioned to allow display of the 2D hand-drawn splashes without visual or mechanical repetition of shape and element.


These 2D flat images, now on distorted cards in 3D space, can be further visually defined with a 3D light to add a greater sense of apparent volume to what is actually a virtual distorted flat surface. These surface images can then be replicated and combined to arrive at visually pleasing compositions of great effects complexity. The controls in the Spryticle system allowed for size, behavior, position, and rendering adjustments that are so necessary in the director-approved production world of frame-by-frame feature animation.

All of this visual complexity was tied to, and directly animated with, a moving 3D water surface animated in Maya. This 3D water surface was built to interact with the 3D location model from the layout department as well as the Golden Gondola that was to be thrashing about, spilling its gigantic pile of 3D golden treasure into the roaring animated 3D water surface.

The visual assembly was a combination of Maya-rendered elements and the broader river water surface, which had been animated in Maya and rendered out using MTOR and shaders specifically written for the water and foam in RenderMan. All of these rendered images were further composited in Animo with the 2D character levels to achieve the final scene images.



As the cistern falls, Spryticles are used to create the effect of splashing water in the climactic ending to DreamWorks Pictures’ animated adventure “The Road to El Dorado.”



A scene from the “Crashing the Gate” sequence demonstrating the use of Spryticle in DreamWorks Pictures’ animated feature “The Road to El Dorado.”

The Use of Traditional and 2D CG Effects as Used to Create the Magic in “The Road to El Dorado”

Jane Gotts
Steve Wood
DreamWorks Animation
jgotts@animdreamworks.com

From the outset, the directors had a very strong idea of what the look of the magic should be. As the magic was primarily to be associated with Tzekel Kan, the evil high priest of “The Road to El Dorado,” it had to have a style and feeling that fit both the character and the Mayan design of the city.

When the creative team got together to discuss how the magic should be created, we realized that we needed to create something unlike anything that had ever been seen in an animated feature film. The tool set we had allowed us to combine 2D traditional effects with innovative 2D CG techniques. This enabled us to use the strength of our drawn design, because we could use a more highly stylized graphic look and combine this with a more abstract ethereal atmosphere that could be achieved with 2D CG.

The overall “look” for the magic came from the color keys of the art director, Ramon Zibach, and the 2D effects animator, Olivier Malric. Olivier’s ideas were fantastic in both senses of the word (dark and mysterious), while still maintaining the strong element of Mayan design essential to the look of the film.


The magic effects ranged from scenes containing a steaming cauldron of green goop to jets of smoke infused with dark and foreboding imagery of demon-like forms. This progressed into mercurial magic

“veins” that flowed over the stone floor, causing cracks to appear and rays of light to shoot through as though from some evil green pit below the earth, which ended by flooding up through Tzekel Kan’s body and propelling him explosively through the air.


The 2D elements were drawn, scanned, painted, etc. and then brought back to effects and the 2D CGI Department. This department has been built up over the last three years from a combination of new tools written for our compositing package, Animo, and a desire to use these tools to push the boundaries of compositing in the color modeling department and the animations produced by the effects department.

The fact that we would receive the scenes in full color gave us an incredible opportunity to add a variety of compositing techniques and a chance to create new 2D CG animation, such as light beams, magic smoke, dust, and animating color changes.

The result, we feel, is that we created a “fresh” look to the magic in “The Road to El Dorado,” something we are constantly trying to achieve in the special effects unit of DreamWorks Animation.



The High Priest Tzekel Kan is about to be transformed in a scene demonstrating 2D CGI light rays, smoke rings, and turbulence composited together in animo for a scene from DreamWorks Pictures’ animated feature “The Road to El Dorado.”



A close-up view of Tzekel Kan performing his magic demonstrates 2D traditional animation combined with various filters composited in Animo for DreamWorks Pictures’ animated adventure “The Road to El Dorado.”

Our developing virtual human can evaluate its body deformation and predict the feeling when it is touched by an object. Current practice in new-product design and development has utilized the motion of a virtual human in evaluating mechanisms for user friendliness and ergonomics. Our focus is to develop a virtual human that can predict feelings such as comfort, pain level, and so on in a virtual space.

When a person uses something, the sensation of touching comes mainly from pressure at a contact surface between the body and the object. We simulate conditions for contact deformation of the virtual human body by a finite-element calculation. Then, we relate the simulated conditions (strain/stress value or load distribution) to the feeling. In this way, we successfully predicted comfortability in using a chair (Figure 1) and pain level in using a bag (Figure 2) by our virtual human.

The prediction of comfortability variation in using a chair at different positions is shown in Figure 1. Comfortability is derived from the load distribution on the contact between buttocks and chair. It is also possible to predict comfortability in using different kind of chairs. Figure 2 demonstrates prediction of pain level in the forearm by contact with the handle of bag. In this case, pain levels are related to strain values in deformation of the forearm.

The virtual human, which we are still developing, will have the following functions and databases:

- Shape data of standard human bodies of men, women, children. Fat and slender factors in their standing positions are also available, and the virtual human can transform itself into their human bodies.

- When a human body changes position, the shape of each part of the body also changes. The virtual human can express the shape conversions of the body that are related to the motion.
- Characteristic data of body material for each part of the human body are available for calculation of contact deformation of the body.
- The virtual human can simulate the conditions of contact deformation of the body by using a finite-element method.
- Data on the relationship between feeling and contact deformation are available, and the virtual human can evaluate the feeling by using the simulated condition of contact deformation.

There are several possible uses for this virtual human, mainly in the medical and welfare fields. It can be applied to design of a nursing-care robot, with special consideration to what level and distribution of forces can be applied to a patient (Figure 3). Another application is designing a bed to prevent back strain. In another application, the simulation can predict women's breast deformation in different positions. This simulation is useful for medical treatment of breast cancer. It can show how the cancer shifts position in different breast shapes.

Note: The figure data of virtual humans used in this sketch are from Poser, MetaCreations Corp.




Figure 1: Prediction of Comfortability in Using Chair (by the result of load distribution on the buttocks)



Figure 2: Prediction of Pain Level in Using Bag (by the result of equivalent strain value on the forearm)



Figure 3: Nursing care robot with special consideration of how forces can be applied to a patient

A Virtual Reality Interface for Analyzing Remotely Sensed Forestry Data

Margaret Watson
Neil Eggleston
Derek Irby
Robert Moorhead
David Evans

Mississippi State University
watson@erc.msstate.edu

Airborne LIDAR (Light Detection and Ranging)¹ has recently become commercially available and shows considerable promise for mapping terrain and land-cover features. Data processing and quality control, however, appear to be lagging behind data-capture abilities. An evaluation of a LIDAR scanning system by the Forest and Wildlife Research Center (FWRC) at Mississippi State University, the USDA Forest Service, EarthData Technologies and two industrial forest companies (Trillium Corporation and Boise Cascade Corporation) inspired development of a CAVE (CAVE Automatic Virtual Environment)² application for exploratory examination and analysis of remotely sensed data from forested landscapes. The primary source of data was an airborne multiple-return LIDAR system flown over 2,000 hectares (5,000 acres) of forest land in coastal Washington.

LIDAR scanning generates files requiring ample disk space and produces datasets that are inherently 3D. Desktop analysis is time-consuming and spatially challenging. Because of the 3D nature of the data and an interest in the relationship between the LIDAR sensor and the actual site, an immersive virtual reality system was considered beneficial for spatial exploration. The Visualization, Analysis, and Imaging Lab in the MSU/NSF Engineering Research Center (ERC) created the CAVE application in collaboration with the FWRC.

In visualizing LIDAR data in the CAVE (Figure 1), the analyst's ability to recognize trends and identify anomalies has been enhanced through various renderings of the data. Spatial interactive tools, such as a compass and a height-measurement tool, have been constructed to orient the forester in the real-world coordinates of the site (Figure 2). Another feature of the measurement tool allows the elevation of

multiple returns to be compared with each other or against reference datasets created from traditional remote-sensing sources and geographic information systems. As a result, forest features such as canopy density, tree height, and crown dimensions are more easily interpreted than was possible using desktop systems.

Visualization of the forest data in the CAVE has also increased understanding of the LIDAR data capture and feature extraction issues. Within the visualization, the data for a class of return can be connected along scan lines. Multiple-return data can be visually linked together to demonstrate each outgoing pulse of the sensor. In the future, on-site measurements of trees and terrain will be introduced into the system to explore the validity of the LIDAR data and the accuracy of analysis algorithms.

ACKNOWLEDGEMENTS

Support for the visualization component of the project began with funding from the ERC and is being continued under a grant from the Remote Sensing Technology Center. The CAVE itself was funded in part by an equipment grant from the National Science Foundation.

www.erc.msstate.edu/~watson/FOREST

References

1. Baltsavias, E. P. (1999). Airborne laser scanning: Existing systems and firms and other resources. *ISPRS Journal of Photogrammetry & Remote Sensing*, 54, 164-198.
2. Cruz-Neira, C., Sandin, D.J., and DeFanti, T.A. (1993). Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Proc. of ACM SIGGRAPH 93*, 135-142.



Figure 1: A small group of individuals examining forest data in the CAVE.




Figure 2: The compass and height measurement tools displaying information about the location of the participant, as well as the height in meters of first (green) and last (blue) return pulses of interpolated LIDAR data.

Visapult: A Prototype Remote and Distributed Visualization Application and Framework

Wes Bethel
Lawrence Berkeley National
Laboratory and R3vis Corporation
ewbethel@lbl.gov

THE PROBLEM

Remote visualization is the process of creating pictures from data, where the data and viewer are on different platforms, often separated by a wide-area network. Traditional approaches to remote visualization have tended to fall into two broad categories:

- Render-remote approaches, where data storage and access, visualization, rendering, and viewer involve a single machine.
- Render-local approaches, which require transmission of data to the viewer's machine, where it is then visualized and rendered.

The former approach has the potential to address large-data problems, but sacrifices interactivity. The latter approach has the potential to preserve interactivity, but assumes that the data will fit entirely on the local workstation, and that it is practical to transmit the data across the network.

OUR APPROACH

We present Visapult, a prototype visualization application and framework that strikes a balance between render-local and render-remote approaches. Visapult was designed to address the needs of scientific researchers who produce and need to visualize hundreds of gigabytes, if not terabytes, of time-varying data. With data of this scale, neither transmission across a network nor local storage is practical, yet interaction with the visualization is crucial. We did not want to sacrifice the gains in understanding that result from motion parallax and stereoscopic presentation.¹ To achieve these goals, we have implemented an application that uses shared rendering. Shared rendering refers to a cooperative visualization and rendering framework in which some of the rendering is performed close to the data, and some rendering is performed on the local workstation. Given the disparity between data transmission rates and the requirements for rendering interactivity, it is desirable to keep frame rates isolated from, and independent of, the latency characteristic of network-based applications.

HOW IT WORKS

Visapult consists of two logical components: a local viewer and a remote, back-end data engine. The back end retrieves scientific data from some source, such as a network data cache,² then performs visualization and partial rendering in parallel. Results are then transmitted to a viewer, which is an interactive 3D graphics application. The viewer and back end communicate over the network using a custom TCP-based protocol. The back-end volume-visualization engine

is an implementation of an image-based rendering assisted-volume renderer.³ The IBR-assisted volume rendering model decomposes nicely into a remote, parallel component and a local, serial component. The remote component scales with available computational resources using an object-order, domain decomposition for parallel-volume rendering. The local viewer, also a parallel application, provides rendering at interactive frame rates, even on low-cost hardware or in stereo. In the viewer, a scene graph model is updated by parallel listener processes whenever new data arrives from the back end, but is rendered continuously by a dedicated process, effectively separating network latency from rendering frame rates. Arbitrary geometry can be included in the scene graph so that volume visualization is complemented with grid visualization, isocontours, or other visualization techniques. The fundamental design paradigm is the combination of parallel network communication, combined with asynchronous scene graph updates.

THE FUTURE


Based upon the results of ongoing performance and profiling analysis, future work in Visapult will focus on performance improvements. From an architectural perspective, aggressive and asynchronous data prefetching will improve overall data throughput. From a "network awareness" perspective, dynamic analysis of the state of the network can lead to parameter changes based on current bit rate, reliability or other quality-of-service attributes. In some cases, bandwidth reservation is needed to allow for visualization of very large (terascale) scientific data, where the scientific data are located on a network cache or are otherwise not local to the back-end visualization and rendering engine.

ACKNOWLEDGMENT


This work was supported by the Director, Office of Science, Office of Basic Energy Science, of the US Department of Energy under Contract No. DE-AC03-76SF00098.

References

1. Ware, C. & Franck, G. (1996). Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics*, 15, (2).
2. Tierney, B., Lee, J., Crowley, B., Holding, M., Hylton, J., & Drake, F.,A. (1999, August). Network-aware distributed storage cache for data intensive environments. *Proc. of IEEE High Performance Distributed Computing Conference*, LBNL-42896 (<http://www.didc.lbl.gov/DPSS/>).
3. Mueller, K., Shareef, N., Huang, J., & Crawfis, R. (1999). BR-assisted volume rendering. In *Proceedings of IEEE Visualization 1999*, Late Breaking Hot Topics.



Raw scientific data are domain-decomposed, then each sub-domain is rendered in parallel by the Visapult back end. The resulting images are transmitted in parallel to the Visapult viewer, where they are inserted into a scene graph as 2D textures.



Visualization of adaptive, hierarchical grids from the combustion simulation are included with volume rendering.

Occlusion culling in z-buffer systems has drawn a lot of attention over the last few years.^{1,3} Due to the high computational costs of occlusion queries, hardware support is required to accelerate interactive rendering. This requires mechanisms that fit nicely into an available graphics subsystem without greedily consuming chip real estate (one reason why hardware support for occlusion culling is rare).

One of the few available dedicated hardware mechanisms for occlusion culling renders objects in an occlusion mode and sets a flag in case parts of the geometry would have altered the frame buffer.² Using a scene hierarchy, bounding primitives can be tested to discard the contained geometry if the flag has not been set. Hence, the bandwidth of the front bus (geometry transfer) and geometry transformation can be significantly reduced, while sacrificing rendering bandwidth during the occlusion test itself. Despite its efficiency, the main drawback of this approach is the costly query, due to latency in the graphics subsystem, which allows only a limited amount of visibility queries performed per frame, and therefore limits the overall possible efficiency.

Visibility-driven rasterization represents a new scheme, which extends rasterization in two ways. First, during frame generation, screen-space-oriented visibility information is generated several times and stored within the rasterizer in a visibility mask. Second, if contributions to the visibility mask were detected, the established visibility information is exploited to discard all triangles and groups of pixels of subsequently rendered polygons that belong to non-visible areas (see Figure 1(a,b)).

Visibility-driven rasterization achieves significantly higher cull-rates than the existing hardware mechanisms. Instead of having a single flag indicating only overall visibility (occlusion culling prior to geometry transformation), we use a frame buffer-oriented visibility mask containing a flag for each tile of pixels to enable culling prior to and during rasterization, in addition to previous culling stages. This tremendously reduces the cost of the unsuccessful visibility queries, by exploiting information during rendering of the geometry, which is at least partially visible. Consequently, more visibility queries can be performed per frame, increasing the overall frame-rate. Visibility-driven rasterization will gain even more importance using multi-texturing, which requires iterating over pixels multiple times, wasting valuable rendering bandwidth while processing occluded pixels.

The performance of visibility-driven rasterization depends on tile size, which again depends on the size of triangles sent down the graphics subsystem. To overcome this dependency, we organize the visibility

mask as register file or on-chip SRAM of 16-bit entries. Hence, a natural tile hierarchy is obtained grouping 4x4 tiles (four 2x2 tiles) into each entry. By using basic bit masking operations (AND/OR), triangles residing across multiple occluded tiles can be easily culled. This reduces the influence of the actual tile size on the culling performance, allowing for a generic tile size selection. We evaluated the quantitative efficiency of visibility-driven rasterization for different tile sizes (extending the MESA 3D graphics library) applied to a set of four real-world scenes: A CAGD cotton picker (10M tris, see Figure 2), a CAGD cathedral (416K tris), a CAGD screwdriver (150K tris), and an extracted iso-surface of the ventricular system of the human brain (270K tris).


On average, best results could be obtained for tiles of 16x16 pixels while the peak culling performance varied from 8x8 to 32x32 pixels. The necessary visibility mask can be stored in as little as 4 Kbit of storage, which can easily be integrated on-chip. For the set of test scenes, this allows for additionally culling between 20 to 40 percent of the remaining triangles prior to rasterization and another 35 to 55 percent of all remaining pixels during rasterization. The latter additional pixel culling is achieved by discarding chunks of pixels of individual triangles that reside across multiple tiles, where at least one tile has been visible (see Figure 1(b)).

One of the advantages of visibility-driven rasterization is its compactness which, in contrast to other approaches, does not require storing depth values (32 bit), since it only contains visibility information. The depth test is inherent to the approach of testing bounding primitives first before rendering the contained geometry. Other approaches might achieve little higher cull-rate, but require major changes to the existing graphics subsystem and consume significantly more chip area to store depth values.

The measured results show that visibility-driven rasterization could extend real-time rendering performance of existing graphics subsystems. Last but not least, visibility-driven rasterization is an efficient culling scheme that fits nicely into the present state-of-the-art, stamp-based rasterization scheme, which is also tile oriented.

References

- Greene, N. (1996). Hierarchical polygon tiling with coverage masks. In *Computer Graphics, Proc. of ACM SIGGRAPH 96*, 65-74.
- Scott, N., Olsen, D., & Gannett, E. (1998, May). An overview of the VISUALIZE fx graphics accelerator hardware. *The Hewlett-Packard Journal*, 28-34.
- Zhang, H., Manocha, D., Hudson, T., & Hoff, K. E. (1997). Visibility culling using hierarchical occlusion maps. In *Computer Graphics, Proc. of ACM SIGGRAPH 97*, 77-88.



Visibility-driven rasterization: (a) Wheel hubs illustrating triangles that can be culled prior to rasterization. (b) Culling of chunks of pixels during rasterization



Cotton-picker dataset containing 10 million triangles.

We have already proposed a volume graphics model for colored pencil drawing (CPD).^{1,2,3} In the model, a regular 3D grid system (volume) is used as the basic data representation, in which the composition elements (voxels) are discretely distinguishable. At the beginning, a paper volume is generated to define the 3D microstructure of paper. Then drawing operations are performed on the volume in order to yield a CPD volume. The drawing operations corresponding to actual CPD techniques are classified into three phases: prearrangement, main work, and coordination. Switching between these phases generates the natural appearance of CPD.

Our previous work includes representative operations such as basic drawing (main work), using a solvent (coordination), and erasing (coordination). These operations are realized as interactions between the texture of the paper and the lead of a pencil or brush. There are a number of additional techniques commonly used in actual CPD,⁴ and direct and faithful reproduction of these techniques is essential in order to represent CPD's charm. In this sketch, we extend our previous model to incorporate three additional CPD techniques: impression, dry wash, and burnishing.

IMPRESSION (PREARRANGEMENT)

The principle of impression is to make a mark (impressed line) that grooves the surface of the paper. After the pencil lead glides gently across the impressed lines, they appear strikingly through the color overlay. Here, we deal with what is called the blind impression only, because it is useful for describing lace or tiny vines in a leaf and is the most frequently used impression technique.

In our model for drawing impressed lines, the thickness of a compressed paper volume is computed according to the impressing pressure. Until the thickness of the paper volume under impressing strokes is reduced to a given threshold, the following procedure is repeated in the vertical direction:

1. An air voxel is plugged up by the upper voxel.
2. The density of a paper voxel is shifted to the lower voxel, and the voxel is changed to an air voxel.

DRY WASH (MAIN WORK)

In dry wash technique, we first use a knife to shave a lead to make fine powder, then rub the powder on a paper surface with a finger or tissue. The technique can produce a fine gradation. Dry wash is utilized for flat shading in a large area.

We assume that a finger with finely powdered lead is rubbed on a paper. The finger area that reaches the paper volume is modeled as an ellipse. Because a finger can not trace the paper texture completely, and the powder is so fine, we can simply assume that the powder adheres to the paper voxels higher than a certain thickness. The thickness is computed according to the pressure at the center point in an ellipse. The powder adheres to hollows in the pulp net whose capacities are larger than the volume of a powder, and to talc voxels at the height field, both of which exist at a higher layer than the threshold. We use a Gaussian-like distribution for describing how the pigment is dispersed away from the two foci of the ellipse.

BURNISHING (COORDINATION)

The general definition of burnishing is to reduce the pigment particles deposited on a paper surface to a fine grain and push them into the surface texture in order to make the surface smooth or shiny. We use an eraser as the burnishing implement. Volume accessibility,⁶ defined by specifying the radius of a spherical probe that is able to reach between objects, is used to evaluate which voxels are accessible to the eraser.


Moving and smoothing pigments in accessible voxels are represented with a 3D extended line integral convolution (LIC).⁶ The concept of accessibility and the principle of LIC are also used in the previous drawing operations.^{1,2,3} The probe radius in the eraser-voxel accessibility determination reflects the softness of an eraser, and the contribution weight for each of the voxels is determined by the pressure and the softness of the eraser and the distance to the target voxel. With a softer eraser and higher pressure, the quantity of removed pigments increases. With a harder eraser and lower pressure, the moving distances of smoothed pigments are longer. The height field after the movement is restricted to a certain threshold, and then the CPD surface is smoothed.

RESULTS

The figure shows a CPD volume that is rendered by a volume visualizer called VolVis.⁷ Impressed lines are used at the contour of the integument of broad beans. Dry wash is applied to the shadowing of the beans and the shading in the background, and burnishing to the luster of some beans. The combination of the present techniques with the previous ones, such as using a solvent and erasing, produces good synergy.

References

1. Takagi, S. & Fujishiro, I. (1997). Microscopic structural modeling of colored pencil drawings. In *SIGGRAPH 97 Visual Proceedings*, 187.
2. Takagi, S., Fujishiro, I., & Nakajima, M. (1999). Volumetric modeling of artistic techniques in colored pencil drawing. In *SIGGRAPH 99 Conference Abstracts & Applications*, 283.
3. Takagi, S., Fujishiro, I., & Nakajima, M. (1999). Volumetric modeling of colored pencil drawing. In *Proc. Pacific Graphics '99*, 250-258.
4. Martin, J. (1997). *The encyclopedia of colored pencil techniques*. Running Press.
5. Miller, G. (1994). Efficient algorithms for local and global accessibility shading. In *Proc. SIGGRAPH 94*, 319-326.
6. Cabral, B. & Leedom, L. C. (1993). Imaging vector fields using line integral convolution. In *Proc. SIGGRAPH 93*, 263-272.
7. Avila, R., et al. (1994). Volvis: A diversified volume visualization system. In *Proc. IEEE Visualization '94*, 31-38.



"Broad beans."

The aim of VScene is to simulate intelligent autonomous behavior by agents in a 3D graphical world, making use of the Unified Modeling Language (UML).¹ Because it is object-oriented, UML is an ideal tool for modeling world items and their interactions. World items comprise objects, including agents, that populate the virtual world. Each autonomous agent is equipped with its own independent perceptual system, consisting of a virtual camera linked to image-processing software.

Learning takes place in a perceive-reason-act loop that takes into account predefined goals implanted into the agent, but each agent is also endowed with the capability of creating new goals. A meta-learning capability using Bayesian inference allows it to reason about the task at hand, to create possible scenarios and to explore alternatives. The Bayesian model uses nodes to represent clauses or events and arcs their causal relations. It was created and developed by Pearl² and has since been used and refined in many contexts, including previous work by one of us (Remagnino).³ Message passing propagates input evidence throughout the net by sending messages between nodes via the links and updating the beliefs of the nodes. An agent has intelligence because, in addition to making decisions, it makes use of learning capabilities to infer the evolution of the scene, and it has autonomy because it infers concepts that are incorporated into its personal world model.

The diagram below shows two such agents. Part of the underlying UML class model for this world is shown to the left of the image and includes classes for the agents themselves, the building and its components,

and the trash cans underneath the steps. On the right is the object model constructed independently by one of the agents. This includes two instances of the "Agent" class, one for itself and the other for the second agent. No object of class "Trashcan" has been instantiated because no object of this class has yet come into the agent's field of view. Although in principle the agent could be endowed with error-free interpretation by being granted access to the underlying VScene object model, the system requires the agent to construct its own model as it interacts with the world. The fact that the agent's model is self-derived allows simulation of increased realistic behavior, including the possibility of mistaken identity and indecision.

Much of the work in this area has concentrated on modeling the behavior of virtual actors. Because object-oriented models capture behavior as well as structure, UML is an ideal tool for modeling all aspects of a virtual world and an agent's dynamic view of it. VScene thus provides a software architecture that allows effective integration of components, including the Bayesian networks, via simple interfaces. This allows greater complexity and robust autonomous behavior in virtual environments.

References

1. Rumbaugh, J., Jacobson, I., & Booch, G. (1999). *The unified modeling language reference manual*. Addison-Wesley Object Technology Series.
2. Pearl, J. (1988). *Probabilistic reasoning intelligent systems: Networks of plausible inference*. New York: Morgan Kaufmann.
3. Remagnino, P., Tan, T., & Baker, K. (1998). Multi-agent visual surveillance of dynamic scenes. *Image and Vision Computing* 16, 529-532.



Water Dome is a hemisphere eight meters in diameter that affords its visitors the unique experience of being surrounded by a water film alive with a variety of projected visuals. We proposed the water dome as an augmented environment.

BACKGROUND

Recent investigations have demonstrated that a wide-angle display such as Omnimax, a spherical screen, might be used to create a virtual immersion experience for its viewers. While fountain technology has been developing for water-screen applications over the past 15 years, projected images have been used in conjunction with mist and streaming water at some theme parks. The water dome is proposed as a method to challenge all five senses in a virtual reality experience. We demonstrated water display,¹ in which one user put his head into the hemispheric water membrane and described the sensation as “being covered with water without getting wet.” The first exhibit was a system for personal use. Next, we created a water dome that many people could enter simultaneously.

THE TECHNICAL ELEMENTS

Falling water creates a display when it strikes an object. We collected data and examined the relationship between the quantity of water to be used and the diameter of the water exhibit. As the flow rate of a pump increases, the diameter of the Water Dome increases as well. A water membrane loses its smoothness when the diameter of the fountain exceeds two meters. It changes into the constituents of streaming water. Regardless of how the water flow is changed, the water display retains a hemisphere that can be entered. In addition, a curtain of water formed by multiple water jets forms a background suitable for image projection.

CONTENTS

The method for projecting images onto the water screen was more important than determining what images to project. Since the Water Dome is a transparent film of water, scattered lights may be seen thorough the Water Dome as participants stand in front of the dome. Furthermore, a projector may be set up in the rear of the dome. Four projectors were placed around the Water Dome so that images on the water surface could be seen from all angles. Because the Water Dome has a curved surface, it provides a distorted screen. In addition, the Water Dome permits double images to be reflected because the dome has front and back layers that lend themselves well to projection of high-contrast geometrical patterns.

IMPLEMENTATION

The Water Dome exhibition project was presented in November 1999 in Tokyo. The flowing water created a liquid dome vivid with color sparkling on the water. Participants inside the dome watched the images projected on the inner side of the dome, while at the same time experiencing a fresh feeling similar to what one feels when standing next to a waterfall. The project provided people with sensations of fun, comfort, and relaxation.

FUTURE DEVELOPMENT


Future water domes could be used by the entertainment industry in amusement parks or as exposition presentations, and they might even find uses in towns as fountains for projection of information.

ACKNOWLEDGEMENT

This exhibition is supported by the Japan Art Scholarship Sixth Grand Prix Award.

Reference

1. Sugihara, Y. & Tachi, S. (1999). Water display. *SIGGRAPH 99 Conference Abstracts and Applications*, 183.



Water Dome visitors view a variety of images while surrounded by water.



The Water Dome exhibit attracted 13,323 people in Tokyo.

The CoMedia (Cooperation Multimedia) project is a joint research initiative between the Fraunhofer IAO in Stuttgart and Ryerson Polytechnic University in Toronto. It is developing a service architecture for delivering media broker and production services. One of the services presently under development is a render service that is made available to various users over the World Wide Web. We can see such a service providing a new way to obtain rendering from competing commercial render entities. Various companies can offer render services to users in a type of “spot market” based on demand, excess capacity, etc. The CoMedia architecture includes an e-Commerce broker to facilitate finding and contracting for all of the services including rendering.

The render engine is: on-demand, parallel, and XML-enabled. On-demand means that the render engine is only invoked when its services are requested. This has implications for cost savings in terms of compute infrastructure. The invocation mechanism is a CORBA-based engine that listens for render requests. Once the render engine is invoked, it determines, via a load-balancing mechanism, available compute resources on accessible networks. Use of CORBA (Common Object Request Broker Architecture) facilitates integration of the rendering service with other CORBA-developed services. This allows a service pipeline to be constructed. For example, after rendering, the rendered frames can be automatically stored in a high-capacity storage service for eventual retrieval. In addition, a CORBA-based video shaper allows the rendered output to be delivered automatically in various forms, such as streamed MPEG for a Web-delivered, real-time test viewing. The final high-resolution files can then be transferred over a high-speed connection or transferred to removable storage media.

In the context of this project, parallel means that the render engine is parallel to its computational core, so calculations such as matrix inversion required by NURBS are performed in parallel. This is accomplished by re-working and expressing algorithms in terms of the MPI (Message Passing Interface for Parallel Computing). This type of parallelism is far more extensive and deep than merely parceling off parts of a frame to be rendered on separate machines and stitching the frames back together. While this approach is more complex, it provides higher performance in terms of the number of frames that can be rendered per second. This approach also allows for the use of multiple low-cost processors. In our test application, for example, we are able to include university computer lab clusters in the parallel render service during off-peak hours. Such an approach is valuable in sharing services. For example, two universities in distant time zones could agree to provide access to enabled rendering clusters to each other during off-peak times. The same approach can be used by other organizations with excess compute capacity at certain time periods. They can elect to make excess services available to a marketplace of users requiring rendering services. The CoMedia architecture facilitates the web-based “marketplace” of buyers and sellers of these services.

XML-enabled means that the render engine is able to read model files that are marked-up using (eXtensible Markup Language). This new approach suggests that XML can be used to provide a universal model mark-up. The intention is that all model packages could export model definitions in a standard XML format, much in the same way that HTML has become a standard exportable file format in most word processors. This is important for providing a generalizable render service that can be applied across product boundaries.

Comedia.rcc.ryerson.ca





Sponsored by ACM SIGGRAPH

ACM SIGGRAPH
1515 Broadway, 17th Floor
New York, New York 10036 USA
+1.212.626.0500
+1.212.764.5537 fax