



VISUALIZATION IN SCIENTIFIC COMPUTING

Computer Graphics • Volume 21 • Number 6 • November 1987
Edited by Bruce H. McCormick, Thomas A. DeFanti, Maxine D. Brown
Published by ACM SIGGRAPH



Computer Graphics is a publication of the Special Interest Group on Computer Graphics of the Association for Computing Machinery. Correspondence to SIGGRAPH may be sent to the chair at the address given below. Materials for editorial consideration should be submitted to the editor-in-chief. Tutorial or other material not generally suitable for journals such as the *ACM*

Transactions on Graphics is of interest. (*Computer Graphics* does not normally publish research contributions.) Announcements, calls for participation, etc. should be sent directly to the production editor with a copy to the editor-in-chief. Letters to the editor will be considered submitted for publication unless otherwise requested. The source of all items published

in *Computer Graphics* will be clearly indicated except for editorial items, where attribution to the editor is omitted. Items attributed to individuals are ordinarily to be construed as personal rather than organizational opinions and in no case does the non-editorial material represent any opinion of the editor as to its quality or correctness.

SIGGRAPH Executive Committee Members

Chair

Kellogg S. Booth
Computer Science Dept.
University of Waterloo
Waterloo, Ontario N2L 3G1 Canada
(519) 888-4534

Vice-Chair for Operations

Bruce Eric Brown
Wang Laboratories Inc.
One Industrial Avenue
Lowell, MA 01851
(617) 967-4297

Vice-Chair for Conference Planning

James J. Thomas
Battelle Northwest
P.O. Box 999 MATH 3000
Richland, WA 99352
(509) 375-2210

Secretary

Beverly Aquino
AHA Management Company
3500 Clayton Rd.
Suite B-185
Concord, CA 94519
(415) 798-7366

Treasurer

Richard L. Phillips
Los Alamos National Laboratory
P.O. Box 1663, MS B-272
Los Alamos, NM 87545
(505) 665-1343

Past Chair

Thomas A. DeFanti
EECS
Univ. of Illinois-Chicago
Box 4348
Chicago, IL 60680
(312) 996-3002

Editor-in-Chief

Richard J. Beach
Xerox PARC
3333 Coyote Hill Rd.
Palo Alto, CA 94304
(415) 494-4822

Directors

Michael J. Bailey
Megatek
9645 Seranton Road
San Diego, CA 92121
(619) 455-5590

Christine A. Barton
Morgan Guaranty Trust Co.
23 Wall Street
New York, NY 10015
(212) 483-4095

Hank Christiansen
Brigham Young University
Dept. of Civil Engineering
368 Clyde Bldg.
Provo, UT 84602
(801) 378-6325

James T. Kajiya
Caltech
256-80 Computer Science
Pasadena, CA 91125
(818) 356-6254

Tom Wright

Computer Associates International
10505 Sorrento Valley Road
San Diego, CA 92121
(619) 452-0170

SIGGRAPH '88

August 1-5 Atlanta

Co-Chairs

Andy Goodrich
University of Michigan Computing Center
1075 Beal Avenue
Ann Arbor, MI 48109
(313) 763-4888

Adele Newton

Computer Science Dept.
University of Waterloo
Waterloo, Ontario N2L 3G1 Canada
(519) 888-4534

Program Chair

John C. Dill
School of Engineering Sciences
Simon Fraser University
Burnaby, B.C. V5A 1S6 Canada
(604) 291-4371

SIGGRAPH '89

July 31-August 4 Boston

Co-Chairs

Branko Gerovac
Digital Equipment Corp.
150 Locke Drive
LM04 H4 Box 1015
Marlboro, MA 01752
(617) 480-6647

Christopher F. Herot

Bitstream
215 First Street
Cambridge, MA 02139
(617) 497-6222

Program Chair

Jeffrey Lane
Digital Equipment Corp.
100 Hamilton Ave.
Palo Alto, CA 94301
(415) 853-6700

Department Editors

References

Baldev Singh
MCC
9390 Research Blvd.
P.O. Box 200195
Austin, TX 78720
(512) 338-3701

Slides

Bruce Eric Brown
Wang Laboratories Inc.
One Industrial Avenue
Lowell, MA 01851
(617) 967-4297

Video

Thomas A. DeFanti
EECS
Univ. of Illinois-Chicago
Box 4348
Chicago, IL 60680
(312) 996-3002

Computer Graphics Cover Editor

Copper Giloth

University of Massachusetts
Art Department
Fine Arts Center
Amherst, MA 01003
(413) 545-1902

Production Editors

Ellen Frisbie/Mary Colton
Smith, Bucklin & Associates, Inc.
111 East Wacker Drive
Chicago, IL 60601
(312) 644-6610

ACM Transactions on Graphics

John C. Beatty

Computer Science Dept.
University of Waterloo
Waterloo, Ontario N2L 3G1 Canada
(519) 888-4534

Eurographics Representative to SIGGRAPH

Carlo Vandoni

CERN Data Handling Division
CH-1211 Geneva 23
Switzerland

Local Groups Coordinator

Stephan R. Keith
424 Custer Road
Hayward, CA 94544
(415) 964-9900
(415) 537-0964

Association for Computing Machinery

11 West 42nd Street
New York, NY 10036
(212) 869-7440

ISSN No. 0097-8930



Visualization in Scientific Computing

Editors

Bruce H. McCormick
Texas A&M University
College Station, TX


Thomas A. DeFanti
University of Illinois at Chicago
Chicago, IL

Maxine D. Brown
University of Illinois at Chicago
Chicago, IL

This material is based upon work supported by the National Science Foundation under Grant ASC-8712231.

Any opinions, findings and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Published July, 1987.



Visualization in Scientific Computing

This report is not copyrighted. Please feel free to use excerpts under the condition that the following bibliographic information is included:

McCormick, B.H. *et al.* (ed),
Visualization in Scientific Computing,
Computer Graphics 21, 6 (November 1987),
ACM SIGGRAPH: New York, p xx-yy.

This report is accompanied by two hours of videotape published as issues 28 and 29 of the **SIGGRAPH Video Review**. See appendix E in this report for descriptions of the contents. Videotapes are available in either VHS or 3/4" U-Matic formats.

To request price information and order copies, contact:

ACM Order Department
PO Box 64145
Baltimore, Maryland 21264
(301) 528-4261

For SIGGRAPH membership information, contact:

ACM
11 West 42nd Street
New York, New York 10036
(212) 869-7440

For SIGGRAPH annual conference information, contact:

SIGGRAPH Conference Office
111 East Wacker Drive
Chicago, Illinois 60601
(312) 644-6610

Acknowledgments

Members of the Panel on Graphics, Image Processing and Workstations

The editors wish to thank the following members of the Panel on Graphics, Image Processing and Workstations for their insightful contributions and comments during the preparation of this report.

David Arnett
University of Chicago
Chicago, IL

James F. Blinn
Jet Propulsion Laboratory
Pasadena, CA

James Clark
Silicon Graphics
Mountain View, CA

Jerome Cox
Washington University
St. Louis, MO

Thomas A. DeFanti
Panel Vice Chair
University of Illinois at Chicago
Chicago, IL

Martin Fischler
SRI International
Menlo Park, CA

Robert Langridge
University of California at San
Francisco
San Francisco, CA

Thomas Lasinski
NASA Ames Research Center
Moffett Field, CA

Patrick Mantey
University of California at Santa
Cruz
Santa Cruz, CA

Bruce H. McCormick
Panel Chair
Texas A&M University
College Station, TX

Mike McGrath
Colorado School of Mines
Golden, CO

Azriel Rosenfeld
University of Maryland
College Park, MD

Alvy Ray Smith
Pixar
San Rafael, CA

Richard Weinberg
University of Southern California
Los Angeles, CA

Turner Whitted
University of North Carolina
Chapel Hill, NC

Karl-Heinz Winkler
Los Alamos National Laboratory
Los Alamos, NM

Visualization in Scientific Computing

Additional contributors

The editors, in an effort to provide accurate and relevant information on a number of highly specialized topics, have relied on the expertise and resources of various industry and university experts. We wish to thank these people for their support and contributions.

Frederick P. Brooks, Jr.
University of North Carolina
Chapel Hill, NC

Carl Machover
Machover Associates Corp.
White Plains, NY

Andrew Hanson
SRI International
Menlo Park, CA

Arthur J. Olson
Research Institute of
Scripps Clinic
La Jolla, CA

Laurin Herr
Pacific Interface
New York, NY

Jeff Posdamer
Washington University
St. Louis, MO

David Hoffman
University of Massachusetts
Amherst, MA

Barna Szabo
Washington University
St. Louis, MO

Members of the NSF Directorate for Computer and Information Science and Engineering, Division of Advanced Scientific Computing

The editors wish to express their appreciation to the NSF Directorate for Computer and Information Science and Engineering (CISE), and its Division of Advanced Scientific Computing (DASC), for ongoing support and encouragement. In particular, we wish to acknowledge the help and assistance of the following individuals.

Gordon Bell
NSF Assistant Director, CISE

John Connolly
Director, DASC

Mel Ciment
Deputy Director, DASC

Albert Harvey
Program Director, New Technologies, DASC

David Salzman
Staff Associate, DASC

Table of Contents

	<u>Page</u>
Title page	i
Acknowledgments	iii
Table of Contents	v
Executive Summary	vii
Preface	ix
Visualization in Scientific Computing (ViSC): Definition, Domain and Recommendations	1
I. Definition of Visualization	3
II. Domain of Visualization	4
III. Recommendations for a ViSC Initiative	8
<u>Appendices</u>	
A. Scientific and Engineering Research Opportunities	A - 1
B. Visualization Environments: Short-term Potential	B - 1
C. Visualization Environments: Long-term Goals	C - 1
D. Industrial Competitiveness	D - 1
E. ViSC Videotape: List of Contributors	E - 1

Panel Report on Visualization in Scientific Computing

Visualization in Scientific Computing (ViSC) is emerging as a major computer-based field, with a body of problems, a commonality of tools and terminology, boundaries, and a cohort of trained personnel. As a tool for applying computers to science, it offers a way to see the unseen. As a technology, Visualization in Scientific Computing promises radical improvements in the human/computer interface and may make human-in-the-loop problems approachable.

Visualization in Scientific Computing can bring enormous leverage to bear on scientific productivity and the potential for major scientific breakthroughs, at a level of influence comparable to that of supercomputers themselves. It can bring advanced methods into technologically intensive industries and promote the effectiveness of the American scientific and engineering communities. Major advances in Visualization in Scientific Computing and effective national diffusion of its technologies will drive techniques for understanding how models evolve computationally, for tightening design cycles, integrating hardware and software tools, and standardizing user interfaces.

Visualization in Scientific Computing will also provide techniques for exploring an important class of computational science problems, relying on cognitive pattern recognition or human-in-the-loop decision making. New methods may include guiding simulations interactively and charting their parameter space graphically in real time. Significantly more complexity can be comprehended through Visualization in Scientific Computing techniques than through classical ones.

The university/industrial research and development cycle is found to be inadequate for Visualization in Scientific Computing. The programs and facilities are not in place for researchers to identify and address problems far enough in advance, even though the emerging

discipline of Visualization in Scientific Computing is found to be critically important to a portion of the country's domestic and export trade threatened by foreign competition. At the present rate of growth, the capabilities of networks, displays, and storage systems will not be adequate for the demands Visualization in Scientific Computing will place on them.

The gigabit bandwidth of the eye/visual cortex system permits much faster perception of geometric and spatial relationships than any other mode, making the power of supercomputers more accessible. Users from industry, universities, medicine and government are largely unable to comprehend or influence the "fire hoses" of data, produced by contemporary sources such as supercomputers and satellites, because of inadequate Visualization in Scientific Computing tools. The current allocation of resources at the national supercomputer centers is considered unbalanced against visualization, in competition with demands for more memory and disks, faster machines, faster networks, and so forth, although all need to be improved.

The Panel recommends a new initiative in Visualization in Scientific Computing, to get visualization tools into "the hands and minds" of scientists. Scientists and engineers would team up with visualization researchers in order to solve graphics, image processing, human/computer interface, or representational problems grounded in the needs and methods of an explicit discipline. The expectation is that visualization tools solving hard, driving problems in one computational science would be portable to problems in another. Proposals would be peer reviewed, and awarded for both facilities and projects at national supercomputer centers and elsewhere. Other agencies of government are encouraged to recognize the value of Visualization in Scientific Computing in their missions and support its development accordingly.

Applying graphics and imaging techniques to computational science is a whole new area of endeavor, which Panel members termed Visualization in Scientific Computing.

Preface

Panel on Graphics, Image Processing and Workstations

In October 1986, the Division of Advanced Scientific Computing (DASC) of the National Science Foundation (NSF) sponsored a meeting of a newly-organized *Panel on Graphics, Image Processing and Workstations* to provide input to DASC on establishing and ordering priorities for acquiring graphics and image processing hardware and software at research institutions doing advanced scientific computing, with particular attention to NSF-funded supercomputer centers. Supercomputer centers had been requesting funds to provide graphics hardware and software to scientific users but, in point of fact, existing tools were not adequate to meet their needs.

Computer graphics and image processing are within *computer science*; the application of computers to the discipline sciences is called *computational science*. Applying graphics and imaging techniques to computational science is a whole new area of endeavor, which Panel members termed *Visualization in Scientific Computing*.

The Panel maintained that visualization in scientific computing is a major emerging computer-based technology warranting significantly enhanced federal support. From the Panel's first meeting came two principal recommendations. It was suggested that the NSF hold a workshop with other government agencies in order to generate a formal summary of the field, and that the NSF establish a new initiative on Visualization in Scientific Computing (ViSC).

Workshop on Visualization in Scientific Computing

The *Workshop on Visualization in Scientific Computing*, held February 9-10, 1987 in the NSF Board Room in Washington D.C., and co-chaired by Panel members Bruce H. McCormick and Thomas A. DeFanti, brought together researchers from academia, industry and government. Computer graphics and computer vision experts analyzed emerging technologies, and federal agency representatives presented their needs and interests. Scientists representing physics, mathematics, chemistry and medical imaging showed examples of their computer-generated imagery using film, videotape and slides. A presentation on Japanese visualization research, a tutorial on state-of-the-art computer graphics animation research, and an overview of commercially available hardware and software rounded out the agenda.

Initiative on Visualization in Scientific Computing

This report presents the findings and recommendations of the Panel for a new initiative in Visualization in Scientific Computing. Much of the impact of visualization, as applied to scientific and engineering research, cannot be conveyed in printed matter alone — so this document is accompanied by a videotape that illustrates pioneering efforts in visualization today.

**Visualization
in
Scientific
Computing (ViSC):**

**Definition, Domain and
Recommendations**

Richard Hamming observed many years ago that "The purpose of [scientific] computing is insight, not numbers." The goal of visualization is to leverage existing scientific methods by providing new scientific insight through visual methods.

ViSC: Definition, Domain and Recommendations

I. Definition of Visualization

Visualization is a method of computing. It transforms the symbolic into the geometric, enabling researchers to *observe* their simulations and computations. Visualization offers a method for seeing the unseen. It enriches the process of scientific discovery and fosters profound and unexpected insights. In many fields it is already revolutionizing the way scientists do science.

Visualization embraces both image understanding and image synthesis. That is, visualization is a tool both for interpreting image data fed into a computer, and for generating images from complex multi-dimensional data sets. It studies those mechanisms in humans and computers which allow them in concert to perceive, use and communicate visual information. Visualization unifies the largely independent but convergent fields of:

- Computer graphics
- Image processing
- Computer vision
- Computer-aided design
- Signal processing
- User interface studies

Richard Hamming observed many years ago that "*The purpose of [scientific] computing is insight, not numbers.*" The goal of visualization is to leverage existing scientific methods by providing new scientific insight through visual methods.

An estimated 50 percent of the brain's neurons are associated with vision. Visualization in scientific computing aims to put that neurological machinery to work.

Today's data sources are such fire hoses of information that all we can do is gather and warehouse the numbers they generate.

Visualization in Scientific Computing

II. Domain of Visualization

Problem:
The *information-without-interpretation* dilemma

Today's data sources are such *fire hoses* of information that all we can do is gather and warehouse the numbers they generate. High-volume data sources include:

- Supercomputers
- Orbiting satellites returning earth resource, military intelligence, weather and astronomical data
- Spacecraft sending planetary and interplanetary data
- Earth-bound radio astronomy arrays
- Instrumental arrays recording geophysical entities, such as ocean temperatures, ocean floor features, tectonic plate and volcanic movements, and seismic reflections from geological strata
- Medical scanners employing various imaging modalities, such as computed transmission and emission tomography, and magnetic resonance imagery

The need to deal with too much data

There is every indication that the number of sources will multiply, as will the data density of these sources. For example, the definition of a supercomputer is changing from its former meaning of 0.1 - 1.0 gigaflops (billions of floating point operations per second) to 1 - 10 gigaflops. Also, current earth resource satellites have resolutions 10-100 times higher than satellites orbited just a few years ago.

Scientists involved in the computational sciences require these data sources in order to conduct significant research. Nonetheless, they are deluged by the flood of data generated. Using an exclusively numerical format, the human brain cannot interpret gigabytes of data each day, so much information now goes to waste.

Earth resource satellites sent up years ago are still transmitting data. All scientists can do is warehouse it. The technology does not exist for receiving, analyzing and presenting information in such a way that it is useful to scientists.

The need to communicate

Visualization by shared communication would be much easier if each of us had a CRT in the forehead. We speak — and for 5000 years have preserved our words. But, we cannot share vision. To this oversight of evolution we owe the retardation of visual communication compared to language. To overcome the

Scientists not only want to analyze data that results from super-computations; they also want to interpret what is happening to the data during super-computations.

ViSC: Definition, Domain and Recommendations

bottleneck, scientists need improved visual interaction (1) with their data and (2) with each other.

1. The ability of medical doctors to diagnose ailments is dependent upon vision. Today, medical imaging is a computational science dependent upon visualization techniques. For example, in hip replacement operations, custom hips can now be fabricated in advance of surgical procedures. Accurate measurements can be made prior to surgery using non-invasive 3D imaging, reducing the number of post-operative body rejections from 30 percent to only 5 percent.

2. Scientists must learn to visually communicate with one another. Much of modern science cannot be expressed in print. DNA sequences, molecular models, medical imaging scans, brain maps, simulated flights through a terrain, simulations of fluid flow, and so on, all need to be communicated visually.

The need to steer calculations

Scientists not only want to analyze data that results from super-computations; they also want to interpret what is happening to the data during super-computations. Researchers want to *steer* calculations in close-to-real-time; they want to be able to change parameters, resolution or representation, and see the effects. They want to drive the scientific discovery process; they want to *interact* with their data.

The most common mode of visualization today at national supercomputer centers is batch. *Batch processing* defines a sequential process: compute, generate images and plots, and then record on paper, videotape or film.

Interactive visual computing is a process whereby scientists communicate with data by manipulating its visual representation during processing. The more sophisticated process of *navigation* allows scientists to *steer*, or dynamically modify, computations while they are occurring. These processes are invaluable tools for scientific discovery.

Immediate visual feedback can help researchers gain insight into scientific processes and anomalies, and can help them discover computational errors. For example, one astrophysicist found an erroneous boundary condition in his code after examining an image of a jet stream with an obvious reflection not apparent in the numbers.

Visualization and Society

The art of visual communication

The English language uses many visual metaphors. We generalize "observation" to mean any perception, call futurists "visionaries", and ask "do you see" when agreement is sought. Certainly, no one can imagine human or mammalian development without sight. Among the professions, the practice of medicine is inconceivable without vision. While eastern cultures revere people skilled in visual communication, such as artists and calligraphers, we westerners take visual skills for granted and tend to hold artists in low esteem.

Mass commercial appeal of visualization

Visualization is a captivating entertainment commodity, as evidenced by society's enthusiasm for video games, rock videos on television, and special effects in feature films. The development of visualization techniques and algorithms for the commercial and entertainment marketplaces, where the objective is to generate realistic-looking images, is already a substantial field of investigation of significant commercial importance.

Scientific potential of visualization

The application of visualization to scientific computing will undoubtedly face a type of cultural inertia well exhibited by the pre-

computer history of visual technology. Over the past 100 years, each newly developed visual medium first mimicked the old. Still cameras were first used to replace landscape and portrait artists. Movies first captured theater from fifth row center; it took 20 years to discover the *vocabulary* to move a camera around. Television neatly adopted the style and content of film; only now are its real-time and interactive capabilities being developed, as in the use of instant replays and graphic overlays. Visualization, the new interactive visual medium, has great potential for new modes of use beyond its origins in rotating logos for television.

Most people see the end result of visualization — reproduced still color photographs or movies. With the exception of flight simulator trainees and video game players, all those not actually in the process of producing visualization see it as one-way and non-interactive. One cannot publish interactive systems in a journal.

The process of scientific discovery, however, is essentially one of error recovery and consequent insight. The most exciting potential of wide-spread availability of visualization tools is not the entrancing movies produced, but the insight gained and the mistakes understood by spotting visual anomalies while computing. Visualization will put the scientist into the computing loop and change the way science is done.

The ability of scientists to visualize complex computations and simulations is absolutely essential to insure the integrity of analyses, to provoke insights and to communicate those insights with others.

ViSC: Definition, Domain and Recommendations

Solution: Visualization technology

Scientists need an alternative to numbers. A technical reality today and a cognitive imperative tomorrow is the use of images. The ability of scientists to visualize complex computations and simulations is absolutely essential to insure the integrity of analyses, to provoke insights and to communicate those insights with others.

Several visually oriented computer-based technologies already exist today. Some have been exploited by the private sector, and off-the-shelf hardware and software can be purchased; others require new developments; and still others open up new research areas. Visualization technology, well integrated into today's workstation, has found practical application in such areas as product design, electronic publishing, media production and manufacturing automation. Management has found that visualization tools make their companies more productive, more competitive and more professional.

So far, however, scientists and academics have been largely untouched by this revolution in computing. Secretaries who prepare manuscripts for scientists have better interactive control and visual feedback with their word processors than scientists have over large computing resources which cost several thousand times as much.

Traditionally, scientific problems that required large-scale computing resources needed all the available computational power to perform the analyses or simulations. The ability to visualize results or guide the calculations themselves requires substantially more computing power. Where will this power come from? The following section considers these needs.

This initiative proposes to place high-quality visualization tools in the hands and minds of research scientists and engineers.

Visualization in Scientific Computing

III. Recommendations for a ViSC Initiative

The ViSC initiative This report supports the implementation of a federally funded ViSC (Visualization in Scientific Computing) initiative. It explores the developing synergy of science and visualization, and makes recommendations to nurture its successful growth and development. This initiative proposes to place high-quality visualization tools in the hands and minds of research scientists and engineers.

Areas of funding The proposed ViSC initiative recommends the funding of both research and technology developments, immediately and in the long-term. Research developments are the responsibility of *tool users*, experts from engineering and the discipline sciences who depend on computations for their research. Technology developments are handled by *tool makers*, the visualization researchers who can develop the necessary hardware, software and systems.

Recommendations for a National ViSC Initiative		
	VISUALIZATION Short-term Needs	VISUALIZATION Long-term Needs
TOOL USERS Computational Scientists and Engineers	Funding to incorporate visualization in current research	Funding to use model visualization environments
TOOL MAKERS Visualization Scientists and Engineers	No funding necessary	Funding to develop model visualization environments

Tool users, short-term visualization needs

Established user communities, funding realities and special needs felt by advanced scientific research facilities, primarily the national supercomputer centers, have encouraged a situation where resources have been devoted to raw megaflops for arithmetic computation, rather than to visualization.

Every scientist and engineer should have a personal workstation...

Every research center should provide on-site capabilities and facilities for high-end visualization.

ViSC: Definition, Domain and Recommendations

Workstations, minicomputers and image computers are significantly more powerful and effective visualization tools than supercomputers. It is a waste of supercomputer cycles to use them to convert model data into new pictures. Specialized graphics processors are more cost-effective than supercomputers for specialized picture processing and/or generation. Workstations should be placed on the desks of each and every researcher to give them immediate access to local graphics capabilities. Every scientist and engineer should have a personal workstation, just as people who must drive for a living need access to cars. Workstations range in price from \$5,000 to \$100,000.

We must also provide concentrated visualization tools to advanced scientific research facilities. Every research center should provide on-site capabilities and facilities for high-end visualization. Visualization equipment and projects should be considered in competition with more memory, more disks, more networking, and so on, to provide a balanced response to user needs.

Current research opportunities requiring visualization are explored in Appendix A. Visualization tools currently available to computational scientists and engineers are described in Appendix B.

Tool users, long-term visualization needs

Commercially available visualization hardware and software are undergoing large-scale development akin to the growth of advanced scientific computing in general. Software is all development cost and no manufacturing cost, so its profitability and maintenance depend on its user base. Hardware does have manufacturing cost, typically 10 to 50 percent of the list price, but this drops with continued sales and the general decline in component prices. Workstations and image computers with impressive power fit on desktops and are well integrated into local networks. Advanced concepts of user interface are being developed and marketed. As long as one's science and engineering needs fit a product line developed for a commercial market niche, all is well.

If researchers only have access to instruments and software supported by pre-existing business demands, they will not be able to push the frontiers of science. Since many new companies are launched as a result of significant scientific discoveries made within research environments, the private sector itself may stagnate if researchers are deprived of adequate tools. Unfortunately, the universities, which have traditionally nurtured basic science without immediate market promise, do not themselves present a large

To encourage the development of visualization tools for scientific and engineering research, interactions must be fostered between scientists, engineers and visualization experts.

Visualization in Scientific Computing

enough market to stimulate the private sector development of leading-edge visualization tools¹.

To encourage the development of visualization tools for scientific and engineering research, interactions must be fostered between scientists, engineers and visualization experts. These interdisciplinary groups should be expected to develop, document, use, and publish both (1) useful results in their discipline, and (2) effective visualization software and documentation. Scientists and engineers need to rely on the experience and intuition of visualization experts to anticipate which representations best convey the information distilled from a cascade of numbers from a supercomputer; the visualization experts need to rely on scientists and engineers to point out the crucial information which flows from the underlying science of a problem.

Hence, we encourage the support of interdisciplinary research teams, rather than just facilities, to ensure that long-term visualization developments be grounded in *real* problems. Close interaction between scientific investigators and visualization technologists will foster better communication between researcher and computer, and between researcher and researcher. The resulting effective and reusable tools can then be shared with scientists and engineers in other research areas, and within the research community at large.

Research opportunities requiring visualization are explored in Appendix A. Needed visualization tools and techniques that must be developed over the next several years are described in Appendix C. America's technological strength in the international marketplace is examined in Appendix D.

Tool makers, short-term visualization needs

At present, commercial industry supports visualization hardware and software. A pressing need is to bring the scientific and engineering communities up to speed using the commercial equipment available. Within this short time span, there is no immediate funding required to develop further tools.

See Appendix B.

1: Academia supports only one product, MOVIE.BYU, distributed by Brigham Young University in Provo, UT, as a maintained set of visualization tools for scientists and engineers.

Visualization and Interdisciplinary Teams

Significant efforts by interdisciplinary teams will produce effective visualization tools. The history of the application of computer graphics to scientific and engineering problems — such as molecular modeling, fluid dynamics or VLSI design — illustrates the importance of close interaction between the application itself and the tools for visualization. Close interaction also makes it much more likely that the tools developed will be reused by other scientists and engineers in other fields, and that their use will then diffuse through their respective communities.

It is expected that all teams engaged in visualization in scientific computing have a mix of skills, and that the development of the tools and techniques of visualization will be an iterative process with different skills contributing at different stages. Here is a list of team members and their associated skills.

Computational scientists¹ and engineers

Computational scientists and engineers conduct research in one of the discipline sciences (molecular modeling, medical imaging, brain structure and function, geosciences, space exploration, and so forth) or engineering (computational fluid dynamics, finite element analysis, and so forth). Given visualization tools, these researchers are responsible for making advances in their areas of expertise.

Visualization scientists and engineers

Visualization scientists and engineers are involved with the development of visualization software, hardware, networks, languages, operating systems and databases. These technologists are developing tools and techniques that have broad applicability to advanced scientific and engineering research.

1: The term *computational science* refers to the use of computers in the discipline sciences, as distinct from *computer science*. The discipline sciences include the problem domains of the physical and life sciences, economics, medicine, much of applied mathematics, and so forth.

Systems support personnel

Systems support personnel understand computer networks, operating systems, programming environments on supercomputers, workstation hardware and software, and device drivers. They integrate systems and new technologies, and configure and maintain visualization facilities for the computational scientist.

Artists

Artists have a formal education in visual communication, including an understanding of color theory, composition, lighting, view angles and related aesthetics. They propose effective visual representations for analysis and communication.

Cognitive scientists

Cognitive scientists, particularly those with formal training in psychology, have experience in human/computer interfaces and visual perception. At the level of basic research, their theories point the field in new conceptual directions. At the task level, they might work with scientists and technologists to develop better user interfaces; for example, they might select input devices or design window management systems.

We believe that future needs will be better served if potential barriers and bottlenecks are confronted now, in a sustained and programmatic way.

Visualization in Scientific Computing

Tool makers, long-term visualization needs

Raw computing power would be more effectively harnessed than it is today if calculations could be understood pictorially and their progress guided dynamically. Modern modes of computing involve interactive, extemporaneous generation of views from masses of data and models, and especially exploration of model spaces by interactive steering of computations. These visual modes offer significantly more usable answers per unit of computing power than symbolic computations alone.

The ability of scientists to comprehend the results of their computations will depend upon the effectiveness of available tools. All users of visualization have an interest in better hardware, software and systems, and much of what is developed can be shared by a number of scientific and engineering disciplines. We believe that future needs will be better served if potential barriers and bottlenecks are confronted now, in a sustained and programmatic way. We believe that the tool makers should be supported to address general visualization issues, such as:

- Interactive steering of simulations and calculations
- Workstation-driven use of supercomputers
- Graphics-oriented programming environments
- Higher-dimensional visualization of scalar, vector and tensor fields
- Dynamic visualization of fields and flows
- High-bandwidth picture networks and protocols
- Massive data-set handling, notably for signal and image processing applications
- Vectorized and parallelized algorithms for graphics and image processing
- Specialized architectures for graphics and image processing
- Establishing a framework for international visualization hardware and software standards

Essentially, funding agencies are being asked to support the development of visualization for the scientific marketplace. Funding would encourage the production of documented, maintained, upward-compatible software and hardware, would guarantee the publication and distribution of results on a variety of media, and would motivate manufacturers to solve bottleneck problems in interfacing and bandwidth.

See Appendices C and D.

To foster a healthy basic research climate with integrated visualization tools, the federal government should commit an amount equivalent to 1 percent of what industry is spending on visualization.

ViSC: Definition, Domain and Recommendations

Funding recommendations

According to Machover Associates Corporation, an international computer graphics consulting firm, the commercial/industrial market will grow from \$7.61B in 1987 to \$23.08B by 1992. A 1985 report by Frost & Sullivan, *Commercial Image Processing Systems in the U.S.*, claims that the world market for image processing (not including the machine vision market) will grow from \$414.8M in 1985 to \$1.58B in 1990, more than a three-fold increase in 5 years.

The Packard-Bromley *Report on the Health of the Universities* calls for doubling the annual federal support for basic research from its present level of about \$7B or \$8B, plus the commitment of another \$5B for badly needed university facilities.

To foster a healthy basic research climate with integrated visualization tools, the federal government should commit an amount equivalent to 1 percent of what industry is spending on visualization.

We believe advanced capabilities for visualization may prove to be as critical as the existence of supercomputers themselves for scientists and engineers.

Visualization in Scientific Computing

Visualization Benefits

This report proposes solutions to an important set of foundational problems. These solutions, if addressed in a methodical, sustained way, offer many benefits to our scientific community at large.

Integrated set of portable tools

Solving problems relies on the efficacy of available tools. Each important client of visualization capabilities, notably industry and mission agencies, has an interest in better hardware, software and systems. We believe that their future needs will be served better if potential technological barriers and bottlenecks are confronted now in a collaborative program.

Scientific progress and leadership

Scientific breakthroughs depend on insight. In our collective experience, better visualization of a problem leads to a better understanding of the underlying science, and often to an appreciation of something profoundly new and unexpected.

Scientific productivity

Better visualization tools would enhance human productivity and improve hardware efficiency. We believe advanced capabilities for visualization may prove to be as critical as the existence of supercomputers themselves for scientists and engineers.

Standardization of research tools

If properly designed and structured, tools and interfaces developed for one discipline science or engineering application would be portable to other projects in other areas.

Safeguard American industrial competitiveness

Our country's technology base, certain mission programs and many high-technology companies will depend increasingly on visualization capabilities. Our industrial competitiveness also has implications for national security. Without a coherent initiative, foreign commodity manufacturers may catch up with American industry from the low end and dominate it.

Making advanced scientific computing facilities useful

The use of today's *advanced* visualization capabilities will eventually spread to industry, medicine and government — beyond the few universities where these capabilities exist today. Supporting a relatively modest but appreciable ViSC initiative now could make supercomputer power usable by these extended communities 5-10 years sooner.

Computational science and engineering encompass a broad range of applications with one common denominator: visualization.

Appendix A. Scientific and Engineering Research Opportunities

Appendix A

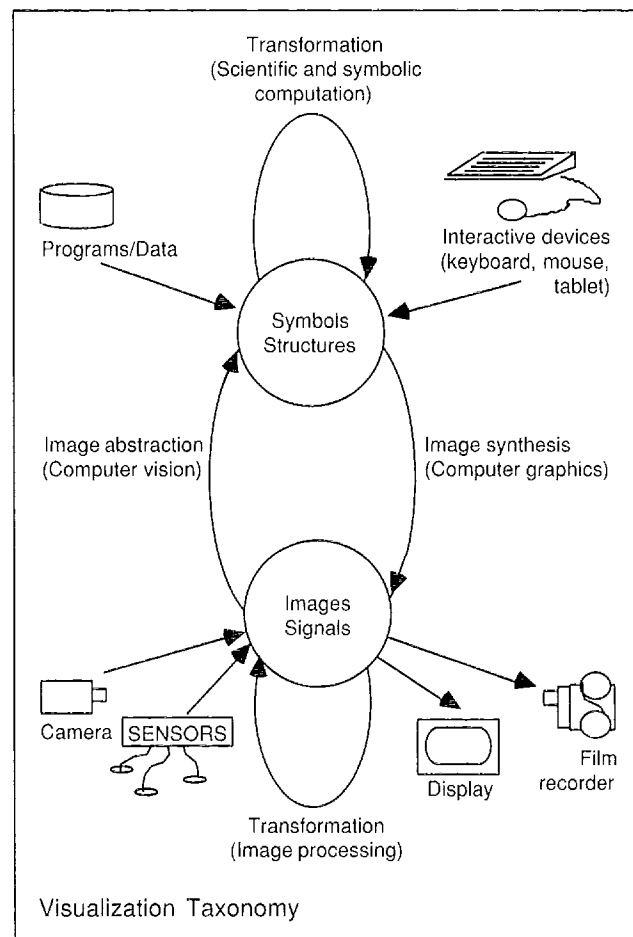
Scientific and Engineering Research Opportunities

Computation is emerging between theory and experiment as a partner in scientific investigation. Computational science and engineering encompass a broad range of applications with one common denominator: visualization. Visualization tools are helping researchers understand and steer computations.

The list of research opportunities for visualization in scientific computing is long and spans all of contemporary scientific endeavor. The research opportunities actually described in this appendix represent a select sampling of advanced scientific and engineering applications, with the constraints imposed by limited space, as well as by the personal biases of the authors.

- *Scientific opportunities:* Molecular modeling, medical imaging, brain structure and function, mathematics, geosciences, space exploration, astrophysics.
- *Engineering opportunities:* Computational fluid dynamics and finite element analysis.

The science of visualization can be defined by the accompanying taxonomy diagram. Images and signals may be captured from cameras or sensors, transformed by image processing, and presented pictorially on hard or soft copy output. Abstractions of these visual representations can be transformed by computer vision to create symbolic representations in the form of symbols and structures. Using computer graphics, symbols or structures can be synthesized into visual representations.



A.1. Research opportunities in science

MOLECULAR MODELING Scientific research opportunities

The use of interactive computer graphics to gain insight into chemical complexity began in 1964. Interactive graphics is now an integral part of academic and industrial research on molecular structures and interactions, and the methodology is being successfully combined with supercomputers to model complex systems such as proteins and DNA. Techniques range from simple black-and-white, bit-mapped representations of small molecules for substructure searches and synthetic analyses, to the most sophisticated 3D color stereographic displays required for advanced work in genetic engineering and drug design.

The attitude of the research and development community toward molecular modeling has changed. What used to be viewed as a sophisticated and expensive way to make pretty pictures for publication is now seen as a valuable tool for the analysis and design of experiments. Molecular graphics complements crystallography, sequencing, chromatography, mass spectrometry, magnetic resonance and the other tools of the experimentalist, and is an experimental tool in its own right. The pharmaceutical industry, especially in the new and flourishing fields of genetic and protein engineering, is increasingly using molecular modeling to design modifications to known drugs, and to propose new therapeutic agents.

Visualization research opportunities

Molecular modeling supports three general activities: synthesis, analysis and communication. Interactive 3D images are essential to each of these areas, to give scientists control of their data and access to information.

Synthesis lets scientists integrate information interactively in real time. The computer is used to build or extend existing models by combining information and knowledge from a variety of sources. Molecular fragments pieced together from a chemical fragment database or a protein structure fitted into a 3D electron density map typifies this modeling activity.

Analysis enables scientists to interpret and evaluate data by selectively displaying experimental and/or computational results in a comprehensible framework. The display and comparison of any

Appendix A. Scientific and Engineering Research Opportunities

number of macromolecular properties, such as chemical composition, connectivity, molecular shape, electrostatic properties, or mobility characteristics, all fall into the domain of modeling analysis. As more structures become available for examination, and as more techniques are developed for analysis, new patterns will emerge. This activity then feeds back to the synthesis activities — and new models for the next level of biomolecular organization can be constructed.

Communication takes place between computer and scientist, and between scientist and scientist. It is important that the information discovered about biological molecules be conveyed not only to the structural scientist, but also to a larger body of scientists whose expertise can add data and knowledge to increase overall understanding. Communication can also bridge the gap between science and the general public, making individuals aware of significant discoveries.

In an effort to make major inroads in these areas, scientists need access to more powerful visualization hardware and software. Effort must be expended on imaginative uses of graphics devices as windows on the microscopic world of the molecule, as well as on integrating the complex numeric and symbolic calculations used to simulate this world.

There are currently two types of images one can generate: *realistic* pictures of molecules (simulations that resemble plastic models), and 3D line drawings (informative images that can be manipulated in real time). Raster equipment is used to create realistic-looking representations and animations, while vector hardware, used for real-time display and interaction, is used to create line drawings. As raster hardware improves, it is expected that raster and vector hardware will merge, allowing increased flexibility in the representations of chemical properties.

Tachistoscopic stereo has been in use for over 15 years and is rapidly gaining acceptance. The introduction of inexpensive liquid polarized screens and polarized glasses will accelerate this important development.

Interaction with the complex 3D world of the molecule is inhibited by the inherent 2D nature of many interactive input devices, such as the mouse. The ability to manipulate the 6 degrees of freedom of a molecule in space so it can interact with another is currently done using dual 3-axis joysticks or similar devices. Imaginative and inventive solutions are needed, such as magnetic motion monitors, to allow multiple interactions.

Visualization in Scientific Computing

MEDICAL IMAGING Scientific research opportunities

Scientific computation applied to medical imaging has created opportunities in diagnostic medicine, surgical planning for orthopedic prostheses, and radiation treatment planning. In each case, these opportunities have been brought about by 2D and 3D visualizations of portions of the body previously inaccessible to view.

Visualization research opportunities

In each of these applications, image processing dominates; both computer vision and computer graphics play a role in orthopedic prostheses and radiation treatment planning. The research activity is experimental, depending on volume-filled images reconstructed from measured data.

The bottleneck in each of these examples is in the generation of useful 3D images. 3D medical imaging requires further visualization research to increase spatial and temporal resolution. Useful 3D visualization algorithms, the development of powerful and portable visualization software, and relevant experimentation in visual psychophysics are all areas of visualization research. The ability to integrate results from multiple modes of imaging seems likely in the future, but has not yet emerged as a trend.

Diagnostic medicine

The imaging modalities of computed transmission and emission tomography, magnetic resonance imaging and ultrasound, enhanced at times by contrast agents or monoclonal antibodies, are leading to a new understanding of both clinical and research questions in diagnosis.

Improved 3D visualization techniques are essential for the comprehension of complex spatial and, in some cases, temporal relationships between anatomical features both within and across imaging modalities. Computation will play an increasingly central role in diagnostic medicine as information is integrated from multiple images and modalities. Visualization, the cornerstone of diagnostic radiology, must be smoothly combined with computation to yield natural and accurate images that a diagnostician can understand in 3D.

Orthopedic prostheses

An emerging visualization application is the fitting of prostheses to individuals for orthopedic reconstructions, such as hip replacements. The 3D fit must be precisely individualized to minimize

Appendix A. Scientific and Engineering Research Opportunities

rejection. Only through non-invasive 3D imaging can accurate specifications be obtained, so that a custom hip replacement can be fabricated in advance of a surgical procedure.

Radiation treatment planning

The use of ionizing radiation to destroy or inhibit the growth of malignant tumors requires careful planning. Misapplication of the radiation beam can jeopardize nearby normal tissue or render therapy ineffective. The precision required for safe but effective treatment is surprisingly high. Fortunately, recent computational advances have made practical the extensive computations that are essential to predict radiation dosage accurately.

Medical imaging allows these predictions to be based on a patient's own anatomy. Before the radiation treatment can be confidently applied, however, an effective means of visualizing the treatment dosage in relation to the tumor and neighboring normal tissue must be developed.

BRAIN STRUCTURE AND FUNCTION

Scientific research opportunities

Visualization in 3D of human brain structure and function is a research frontier of far-reaching importance. The complexity of the brain limits understanding gained from the purely reductionistic approach familiar to neurobiologists. For continued progress in brain research, it will be necessary to integrate structural and functional information at many levels of abstraction.

Work on brain structure and function requires computational support in four areas:

1. Acquisition of experimental data in digital image form from serial histological sections, medical imaging instruments, drug receptor studies and neurophysiological experiments.
2. Extraction of features from measured digital images to produce a 3D map of brain structure and function.
3. Analysis of the abstract brain map to relate measured images and parameters to a standard brain geometry, to provide statistical summaries across a series of brains and to compare an individual brain with such statistical summaries.
4. Visualization of the results of data acquisition, feature extraction and map analysis in a proper 3D context.

Visualization research opportunities

The massive data input needed to map the brain will eventually lead to the invention or use of more advanced storage technologies. However, there is also a gap in our present ability to do automatic feature extraction. Manual recognition of features is inadequate for so large a problem. Automation will require major advances in the field of volume image abstraction and modeling.

While the accomplishment of these tasks remains a monumental challenge, selected pilot studies promise near-term resolution of several pivotal issues. Brain mapping is a necessary first step toward modeling and simulating biological brain functions at a systems level of description.

3D brain visualization

Large-scale volume memory modules capable of storing 1 gigabyte of volume data are under design to model the brain from serial histological sections. Renderings are then constructed from the 512x512x512 array of voxels to compare the spatial distribution of neurons and brain structures within one brain or, at increased resolution, within one portion of the brain. Image processing can create complex representations with brain parts made transparent, translucent or opaque. The volume memory module will aid in the use and evaluation of spatial filtering and boundary detection techniques in 3D brain imagery.

3D image understanding

A neuroanatomist locates objects in a brain section by matching the contour and regional data with a memorized model drawn from a visual knowledge base, which can include references to atlases. An *image understanding system* provides similar computer-based guidance for the semi-automated image analysis task. For example, at the stage of data input, an individual can first identify a feature in one section and then let the system track the feature in succeeding serial sections, label the sectional data and store the information about the feature in a hierarchical database. Mapping of brain interconnections in parallel throughout the brain, at least at the level of nerve fiber tracts, remains a shortcoming in our mapping technology.

Brain mapping factory

Useful statistical experiments will eventually require analysis of 50-100 brains or major brain portions. The magnitude of this task is such that special automated facilities, called brain mapping factories, will inevitably be required to map one brain per month. Super image computers of the computational power required (1 instruction per

Appendix A. Scientific and Engineering Research Opportunities

nanosecond) will be available within a few years. By centralizing the volume image analysis at a brain mapping factory, it becomes possible to integrate and standardize scanning instrumentation, stain technology, image analysis software, geometric modeling and other brain mapping technologies, and to exploit economies of scale. Subsequent brain analysis could take place at distributed workstations with enhanced tools for image analysis.

MATHEMATICS Scientific research opportunities

In the computational study of partial differential equations associated with gas dynamics, vortex formation, combustion and fluid flow, the most effective means of analyzing output has been visualization. Using modern supercomputers, novel parallel architectures and new mathematical algorithms, important 3D physical processes can now be simulated. Geometric problems, such as the generation of body-centered coordinates, automatic mesh generation, and so on, are mathematical in nature and will have to be solved efficiently to conduct research in applied mathematics.

Visualization research opportunities

Visualization is making a tremendous impact in the mathematical study of optimal form and more generally the calculus of variations, including the theory of minimal surfaces, and surfaces of constant mean curvature. Computer graphics has become an essential research tool in this and many other areas of pure and applied mathematics. Hard problems, such as eigenvalue optimization for regions with partially-free boundaries, are being attacked successfully for the first time with these visualization tools.

Mathematics is one of the last sciences to become *computerized*; yet, it is already clear that visualization, coupled with very high-speed numerical simulations, is having a major influence in the field, even in areas long considered to be abstract. This new mode of investigation makes collaboration with scientists in other disciplines much easier for the mathematician; there is a common language of computation and images. There is a strong need to increase the availability and power of visualization tools for researchers within this discipline.

GEOSCIENCES Scientific research opportunities

Generating and coordinating symbolic descriptions of complex sensor data is an extremely challenging research problem. A variety of such data are collected as source material for analysis in the geosciences, including photographs, long and short wavelength imagery, synthetic aperture radar, seismic sensor signals, and topographic and oceanographic surveys.

When analyzing geoscientific data using automated methods, one constructs models of terrain and cultural features, such as elevation, bodies of water, drainage patterns, vegetation, land use, mineral deposits, roads, buildings, urban areas, and so on.

Visualization research opportunities

Many different methods may be used to generate relevant information from sensor data, but virtually all involve transforming the data into a variety of graphic displays. Highly capable techniques for visually representing the information content of such data sources will become critical research tools in the near future.

In all the geoscientific fields, the trend is towards using more powerful computer tools to convert the analysis of sensor data into visual physical structures that can be inspected interactively by a scientist. The combination of material models and sensor data also permits simulation of physical processes for which actual observations do not exist. Another important direction is the use of parallel hardware to process massive data sets.

Testing the validity of new scientific concepts requires another form of visualization. Scientists want to be able to view intermediate results in real time to determine whether their simulation programs are correct.

Cartography

In the cartographic compilation process, raw data is analyzed, reconciled and assembled into cartographically useful symbolic representations. Manual compilation is time-consuming, tedious and has few tools for referring quickly to previously compiled information. The variety of views and displays of specialized information required by users for interpretive purposes cannot be supplied by conventional maps.

Compilation process: With visualization techniques, an analyst can work with many different, related data sets on a display screen; for

Appendix A. Scientific and Engineering Research Opportunities

example, one can use computer-generated markers to show exactly how different images correspond.

Synthetic information displays: Automated synthetic information displays supply insights not available with manual techniques. Not only is an analyst's efficiency improved, but special-purpose cartographic displays can be constructed and tailored to a user's needs in lieu of conventional maps. Visualization systems extend the user's ability to access compiled information by providing data, such as simulated stereographic images, showing the natural appearance of a geographic area with specific feature overlays. In addition to providing static images, such systems can generate simulated flights through a terrain, giving a viewer an immeasurably better intuitive understanding of a geographic environment than is possible with paper maps.

Meteorology

Meteorological predictions depend upon an analyst's ability to understand and visualize time-varying flow patterns in the atmosphere. (See Computational Fluid Dynamics.)

Geology

The interpretation of information such as seismic data and well logs is a critical part of the process of locating mineral deposits. The original data needs to undergo a great deal of processing in order to derive its topographic meaning.

Visualization tools are needed to detect significant trends in the original data, to understand the mapping of the original data to geographic locations, and to do trial-and-error analysis using a variety of options until results consistent with known external information are obtained.

SPACE EXPLORATION

Scientific research opportunities

The field of planetary study is accumulating huge volumes of data on the planets in the solar system. The issues cited for geology and meteorology in the Geosciences section are valid, whether talking about Earth or the rest of the solar system. The section on Astrophysics also describes research of significance to this discipline. Enough data is now available that scientists are beginning to integrate observed phenomena and theory from the fields involved in planetary study: meteorology, geography, planetary physics, astronomy and astrophysics.

Visualization in Scientific Computing

For the last several decades, scientists' abilities to gather data has surpassed their abilities to analyze and digest it. Access to supercomputers and visualization tools will benefit a number of areas:

1. The integration of multi-dimensional databases for 3D reconstruction of planetary, or stellar, environments.
2. The simulation of supernovas and the associated analysis of new data to determine thermal dynamics, fluid dynamics and nuclear reactions inherent in the event.
3. The study of convection cells in solar interiors.
4. The study of solar seismology.
5. The observation of doppler shifts over the surface of the sun, spatially and temporally, to infer acoustic modulation and determine vertical and horizontal velocities in order to resolve internal and surface solar structure.
6. The determination of density functions of galaxies.
7. The ability to look at colliding galaxies as a many-bodies problem and to observe tidal effects of proximate galaxies.
8. The analysis of magnetic field dynamics and associated effects.
9. The ability to model and study geodynamics in planets and satellites.
10. The hyperspectral analysis in geology and biology.
11. The ability to search and correlate multi-channel radio frequency data for the Search for ExtraTerrestrial Intelligence (SETI).

Visualization research opportunities

The problem of correlating multiple, large databases has the parallel problem of presenting the complex databases generated. The ability to communicate phenomena that are not only visually difficult to portray, but are also non-visible, multi-spectral or multi-channel, requires visualization research. Regardless of whether the data is displayed as a still image or an animation — in 2D or 3D, realistically or schematically — the compactness of the data in a well-designed image is of great benefit to the scientific community.

Visualization is an especially effective tool for communication among scientists in different fields of space exploration. All these fields require a format for the distillation of huge amounts of numerical data. Access to visualization is critical. These tools do not simply represent the best way to look at data; in fact they represent the only way to see what is going on.

Appendix A. Scientific and Engineering Research Opportunities

ASTROPHYSICS **Scientific research opportunities**

From Galileo's drawings of sunspots to images of extragalactic radio jets synthesized at the Very Large Array (VLA) on the plains of New Mexico, observational astronomy has been a science with a preeminently visual orientation. With growing access to supercomputers, the computational scientist now joins the theorist and experimentalist in facing the challenge of assembling vast sets of data in intellectually digestible form. Whether the data is obtained by observations with the most advanced imaging detectors or by simulations with the most sophisticated computers, the problem is one of visualization.

Astronomical events and objects can now be simulated with a detail and complexity approaching that of the observations. For example, the theory of supernova explosions, developed on supercomputers of the 1960's, correctly predicted the nature of the neutrino burst detected from the supernova 1987a in the Large Cloud of Magellan. With vastly improved algorithms for fluid dynamics, general relativistic dynamics, thermonuclear reaction networks, and radiation flow, supercomputers of the 1980's allow scientists to simulate the complex behavior of matter accreting onto neutron stars, the birth of galaxies after the Big Bang, the seismology of our sun, or the formation of a planetary system. Problems so difficult that they have resisted attack for years, such as the study of stellar (and solar) convection, are beginning to be reexamined with these tools.

The general theme is well represented by astrophysics, but is much broader. With increased computer power, the computational scientist, the theorist and the experimentalist can be more ambitious. The area of interaction between these three elements of science is visualization.

Visualization research opportunities

An obvious and enduring problem in modern astrophysics is data access. When astronomers kept data on photographic plates, the solution to problems was obvious: keep the plates and look at them. Indexing and retrieving data were difficult because there were a lot of astronomers and a lot of plates. With the many new channels of information now available, the problem becomes deeper and more acute. Also, some modes of information do not have unique visualizations, such as multi-frequency radio maps or cosmic ray energy and abundance spectra.

Visualization in Scientific Computing

If the primary benefit from an experimental program is a mental image that summarizes the data, then the importance of visualization becomes clear. The major bottleneck to successful and efficient use of supercomputers is the human/computer interface; that is, scientists need a mechanism for converting literally millions of numbers into ideas, and the ability to interactively, graphically interrogate data sets in which those ideas supposedly reside.

Appendix A. Scientific and Engineering Research Opportunities

A.2. Research opportunities in engineering

COMPUTATIONAL FLUID DYNAMICS (CFD)

Scientific research opportunities

CFD refers to the numerical solution of the partial differential equations of fluid flow, the Navier-Stokes equations. These equations are central to fields of applied research study such as aeronautics, automotive design, weather forecasting and oceanography, and also to more fundamental research in the understanding of fluid dynamics.

Visualization research opportunities

Within the field of CFD, issues of scientific computing (solution algorithms) and visualization (techniques for viewing complex vector fields) are predominant. CFD is theoretical in nature with detailed simulation of fluid flows its principal goal. The accuracy and complexity of CFD simulations have grown significantly with the recent availability of supercomputers. As a consequence, 3D Navier-Stokes flow computations for complicated geometries are now feasible. Visualization of such calculations is essential to understand the wealth of information in such simulations.

While the main bottleneck to progress is visualization of flow field simulations, the organization of large amounts of input data and the complexity of 3D vector fields are also challenging aspects of the CFD visualization problem. It is anticipated that CFD research will increasingly rely on visualization techniques to allow the CFD scientists to absorb and understand the results of supercomputer simulations.

Identification of flow structures

There are many known structures associated with fluid flows, such as shock waves, vortices, shear layers and wakes. Software algorithms and visualization techniques are needed to identify these structures in 3D flow field solutions. In areas such as turbulence research, the fluid structures are often unknown. Visualization techniques to help discover and define new structures in turbulent flow would be particularly valuable to this area of research.

Display of CFD data

Innovative techniques for displaying 3D data are required both in hardware and software. Techniques for data transformation, hidden surface removal, transparency and multi-dimensional cues should normally be available in near real time (5-10 frames per second). This

Visualization in Scientific Computing

would provide an effective interactive environment in which the fluids researcher could analyze numerical simulations.

Surface geometry and grid generation

The accurate solution of the fluid equations requires good numerical definitions of the body over which the fluid flows. Software is required to expedite the building of 3D models. The proper discretization of the space around the body, such as building a 3D grid, is often critical to obtaining a valid numerical solution. The problem of how to build and visualize such grids over complex 3D objects is a central area of research in CFD.

Comparison of experimental and computed flow fields

To validate CFD computer programs requires the comparison of simulations with experimental results. Image processing techniques are required to analyze videotape or digitally-stored images. The actual comparison of computed and experimental flow fields through direct visualization techniques is virtually unexplored in CFD.

FINITE ELEMENT ANALYSIS

Scientific research opportunities

Finite element analysis is a computational technique for the solution of a variety of problems in such diverse fields as aircraft design, dam construction, modeling the motion of the inner ear, analysis of the stress in building components, and tectonic plate motion. In fact, finite element analysis is the foundation of the solution phase of modern computer-aided design (CAD).

From a mathematical point of view, finite element methods are a family of techniques for the approximate solution of partial differential equations through one of several variational formulations. These methods are characterized by the fact that the approximating functions are constructed from functions defined on a few simple subdomains called finite elements. Since 1956, this method has gained general acceptance in all branches of engineering and computational physics.

The main goal of finite element research today is to develop procedures to assure the quality and reliability of many kinds of engineering and scientific computations involving field equations. Research is active in the following areas: implementation of solid models; treatment of nonlinear phenomena, especially problems of dynamics and stability; and adaptive selection of an analysis mesh.

Appendix A. Scientific and Engineering Research Opportunities

Visualization research opportunities

Visualization techniques play very important roles in both the industrial application of the finite element method and in research directed at improving the reliability and accuracy of the method.

Solid modeling started in the late 1960's using Boolean set operators on simple solid primitives. Integration with finite element meshing and post-processing displays are of more recent origin. At present, little research is being done to integrate solid modeling with the latest thinking in the design of hierarchical finite element spaces. A number of problems of great practical importance can only be solved by proper integration of solid modeling, optimal meshing and artificial intelligence techniques. The development of advanced 3D imaging is an essential part of technology development in this area.

Solid modeling

Understanding the basic design principles of 2D finite element meshes is quite recent (1984). Proper mesh design, coupled with p-extension, now makes it possible to realize nearly optimal (exponential) rates of convergence. Very little, however, is known about optimal mesh design for 3D objects. Integration of finite element technology with solid modeling requires better understanding of mesh design procedures and what kinds of surface mapping techniques should be incorporated into the finite element modules.

Nonlinear phenomena

Nonlinear phenomena frequently appears in the study of dynamics and stability of a body or field. Correct visualization techniques are essential for the understanding and numerical treatment of these important physical phenomena. This field requires techniques for the visualization of solution manifolds.

Adaptive model selection

There are substantial complexities associated with the development of models for engineering decision-making which involve the selection of the proper idealization and the creation of a mesh of this idealization. These decisions depend upon the goals of the computation and the guidance derived from the application of artificial intelligence techniques.

This research involves the visual representation of a finite element solution in a form suitable for a user to conveniently evaluate the results and, using the information generated in the analysis cycle, interact with intelligent guidance systems. This extremely important field will drive the field of computer-aided design.

A.3. Visualization research tools: the common denominators

All of the research opportunities just described are characterized by an urgent need for visualization. Computer graphics, computer vision and image processing are all required in varying degrees. In every case, limited visualization capabilities and limited access to visualization facilities are major bottlenecks to progress.

More specifically, we can pinpoint one visualization paradigm which is growing in popularity, but which will require substantial research: *interactivity*, a tool that enables scientists to steer computations. It will enrich the process of scientific discovery, lead to new algorithms, and change the way scientists do science.

Interactivity and 3D surface models

Animation of complex objects and scenes will be of increasing importance in providing scientific insight. Animated 3D surface depictions will be used for most of molecular modeling, orthopedic prostheses in medical imaging, surface models of brain parts and descriptions of a body over which there is fluid flow. Animated surface renderings will also provide the main currency for the description of robot manipulation, the design of manufactured parts and for the layered depositions of microelectronic chips. There are several trends taking place in animation:

1. Imagery of increased complexity needs to be visualized within reasonable time frames.
2. Highly articulated structures require the simulation of a *nervous system* of their own to coordinate and control behavior that is becoming increasingly complicated.
3. The ability of a user to navigate through a 3D scene is needed.
4. Feedback permits the underlying assumptions of a model to be perturbed.

Interactivity and volume visualization

The central importance of volume visualization in scientific computing is largely unrecognized because of people's long experience with surface-rendered images. This field is very new, and there is much research to be done to provide the necessary tools and techniques. The most difficult, perhaps, is providing the aspect of interactivity.

Visualization will increasingly require navigation in volume-filled imagery. Additional complications will come from the portrayal of vector and tensor fields in 3D volumes. In the future, the animated and interactive presentation and analysis of volume imagery will play a

Appendix A. Scientific and Engineering Research Opportunities

central role in medical imaging, brain mapping, geosciences, astrophysics, computational fluid dynamics and finite element analysis, as well as many other scientific and engineering fields of research.

A glimpse of the future: Using visualization tools to solve problems

Consider the following application of the future. An airplane designer wants to design an airfoil and test it by simulating air flows over it. There are two ways to represent this problem: surfaces or volumes. Using a surface description, the airfoil is rendered as an optically accurate representation of the metal or ceramic mechanical parts from which it will be made; it will have the correct reflectances, colors, opacities and surface attributes. Using a volume description, the volume of the airfoil is displayed. Instead of viewing its optical characteristics, the engineer selects an attribute to observe at every point (voxel). Assume stress is selected; the engineer then assigns a color to this attribute — for example, gray could be selected to represent low values of stress, red for high values, and a color from the ramp between these two extremes for intermediate values. (Note that the designer could have used the surface device to view surface stress. For the purposes of this example, however, the surface representation is used to generate realistic-looking imagery, and the volume representation is used to display the distribution of stress within the airfoil.)

Further, the air flow around the foil represents air turbulence around the wing. Turbulence is represented as vector fields on the surface device. On the volume device, flow particles vertically near the wing are portrayed in a bright color, such as green, while particles further away are darker and more transparent.

Using animation, the engineer can see the shape (external and/or internal) of the airfoil, the turbulence in the boundary layer air flowing over it, and, in the volume image, the stresses in the airfoil. The designer can interact with either display in two ways: *interactively* changing the angle of the 3D display in real time or *steering* the simulation. The designer can change the shape of the wing in real time, the speed at which the wing is flying through the air, or the altitude and hence the characteristics of the air.

The designer simulates the new design in real time, alters it in real time, and steers a simulation — all by using visualization techniques. The engineer also saves millions of dollars by not building early prototypes to check the integrity of the design.

Appendix B

Visualization Environments: Short-term Potential

Visualization is an emerging technology that enables better computational science. Scientists already generate such torrents of information that they can no longer do much more than gather and warehouse it. Visualizing the results of complex computations and simulations is absolutely essential to insure the integrity of the analyses, to provoke insights and to communicate them to others.

In time, access to visualization tools will harness new sources of power for exploration, discovery and communication. But, what is feasible today? What visualization tools are currently in use? This appendix examines the current and emerging visualization environments available to scientists today:

- Visualization data types and software
- Visualization equipment
- Visualization facilities
- Network usage

B.1. Visualization Data Types and Software

DATA TYPES There are two types of spatial data representations in visualization: two-dimensional and multi-dimensional (the most common being three-dimensional).

2D data The most common form of visualization today is the 2D chart or graph. We think of 2D pictures as having no depth information; they are often plots of two variables, but can take advantage of gray scales and color to represent additional information. For example, a 2D chart can apply pseudo-colors to array data to represent a third parameter value, such as density.

2D image data is often stored as either data arrays or *pixel* (picture element) arrays, such as bit or byte maps, without any underlying structural descriptions. The size of the array is dependent upon the resolution of the image.

**3D data
(multi-dimensional)**

3D data can be represented as surface or volume models. We are much more familiar with the former, since 3D points and vectors date from the early 1960's, and 3D surfaces, commonly used in CAD/CAM, date from the early 1970's. *Voxels* (volume elements) as 3D generalizations of 2D pixels began to appear in the literature in the mid-70's, particularly in medical imaging, and voxel machines have been built. In the late 1980's, hardware and software imaging technology is advancing so that we can represent the actual contents of a volume, not just its surface.

3D surface models are stored as programs and data rather than as pixel maps to preserve underlying structure descriptions and to economize on storage space.

3D surface models: complexity measures

Surface rendering: Since the resolution of an output medium changes with the application, the time it takes a program to build objects and render them for display with surface attributes is the complexity measure of choice. Rendering varies from real time (usually with considerable hardware assistance and cost) to minutes or hours per frame. Surface models can be built and rendered using supercomputers and/or workstations; tradeoffs take into consideration the cost and speed of the hardware, and the complexity and resolution of the model to be generated.

Appendix B. Visualization Environments: Short-term Potential

Supercomputers are required when rendering software is not supported by workstation firmware or when pushing the frontiers of the number of floating-point operations carried out per pixel.

Storage capacity: Uncompressed 2D images of 3D data can vary from 256 kilobytes (512x512 with 256 colors) to 24 megabytes (2048x2048 with 2^{36} colors plus 12 bits of mixing information); this represents a difference of 2 orders of magnitude. Compressions of an image often can reduce storage capacity by 90 percent. If the image is generated synthetically rather than captured from nature, then storage requirements can be considerably reduced by storing only the program and the model data.

3D volume models: complexity measures

Data computation and transmission: Voxel data, such as interpolated sectional data from medical imaging scanners or seismologic sensors, presents staggering needs for data computation and transmission. Computing the color value of each voxel of a 512x512x512 3D image takes over 134 million memory cycles, even at one memory cycle per voxel. Processing requirements often exceed 1000 cycles/voxel, a significant amount of time per image even on a dedicated supercomputer.

Storage capacity: Current volume displays hold 256x256x256 voxels. In the near future, displays will increase to 1024x1024x1024, requiring storage capacities 2 orders of magnitude greater than what is currently available.

2D projections of multi-dimensional data sets

Many sensors collect data that are functions of more than two parameters. To visualize multi-dimensional data using conventional 2D displays, various methods of selecting 2D subspaces of the data can be used. For example, one can display planar cross sections of the data, or projections of the data on a user-selected plane. A more interesting approach, however, is to automatically select 2D subspaces of a high-dimensional space for display; for example, the computer might choose a plane defined by the two eigenvectors of the data that have the largest eigenvalues. Non-linear subspaces of the data, such as curved surfaces rather than planes, could also be used. Techniques for automatically selecting 2D subspaces of higher-dimensional spaces have been extensively studied by the developers of interactive pattern recognition systems.

Color displays are very useful for displaying vector data of dimension less than or equal to three, collected by sensors or computationally generated. When dealing with higher-dimensional vectors, such as the multi-spectral bands collected by earth resource sensors, one can display lower dimensional subspaces using color coordinates to represent any three selected bands. A more interesting idea is to compute up to three scalar-valued functions of the data and display these function values as color coordinates. For example, one can display selected linear combinations of the bands, where the coefficients are chosen to correspond to the eigenvectors that have the largest eigenvalues. The combinations used need not be linear; for example, ratios of bands often provide useful information. Techniques like these have been extensively studied by the developers of systems for analyzing and displaying multi-spectral remote sensing data.

One often wants to display not the original data, but new data derived from it using various image processing techniques. A wide variety of operations are described in image processing and computer vision, such as:

- Linear transforms, such as the Fourier transform
- Geometric transforms, such as rotating, rescaling, projecting, warping, and so on
- Local averages of the results of smoothing operations, such as medium filtering
- Derivatives, for edge enhancements or extractions
- Feature maps, showing various types of local features in the data
- Results of various segmentation processes applied to the data

3D DATA MODELING

This section addresses 3D surface models, since volume graphics and imaging are still in their infancy. Volume visualization will play a major role in future research environments, as described in Appendices A and C.

A scientist creates a model to help interpret, evaluate and present an analysis. The investigator can also use the resulting images to refine or redirect a simulation, either in batch mode or in real time, depending on what hardware capabilities are supported.

The visualization software currently available at scientific facilities can be categorized into four functional areas: defining models, constructing models, rendering models and displaying models.

Appendix B. Visualization Environments: Short-term Potential

Defining a model Model definition is the process of creating an object database. Data is either defined by geometric primitives in a graphical database, or by schemas in a visual knowledge base.

Constructing a model: simulation and/or analysis Construction is the computationally intensive task of building a model; this component of visualization is where the power of supercomputers is most needed. If the data is defined by points and lines, then computer programs build, or *simulate*, the model described. If the data is an image, then computer programs *analyze* the input and extract new images or numeric data. For example, an engineer might want to simulate the air flow over an aircraft wing in a wind tunnel, or analyze the structural load on a new automobile part.

Rendering a model Regardless of whether the original image came from a geometric database or a visual knowledge base, the rendering software takes the model structure and produces images, usually with color, hidden surfaces removed, lighting sources and surface attributes. Rendering is computationally intensive, and can be done either on supercomputers or specialized graphics workstations.

Displaying a model Display software is the *glue* used to combine or enhance images generated or analyzed by the components above, such as compositing tools, fonts and text, picture file saving/restoring, retouching utilities, color manipulation, and so on. It also consists of device drivers that deal with the idiosyncrasies of I/O devices; device drivers accept user input and present computer-generated images on display hardware for user viewing or recording.

Visualization Software

LINES

The earliest software for graphics drew lines in 3D and projected them onto a 2D plane, offered viewing transformations for looking at the result, and offered transformations (scale, rotate, and translate) for describing the line objects. The theory and practice of drawing lines, expressed in homogeneous coordinates, and the control and display of lines using 4x4 matrices, represented a major development in computer graphics in its time. It is described in many texts and is built into hardware vector generators currently available.

A variety of standards in use today incorporates these basic principles. The CAD/CAM industry has embraced this level of the art. It is cheap enough to put on every engineer's desk and fast enough for real-time interaction.

The representation of surfaces by grids of lines is well-known and is often used to signify to the public that something is *computer-generated*. Splines, polygons and patches are now commonly provided in line drawing software packages, although there is still insufficient commonality between packages as to which types of splines and which types of patches are provided.

POLYGONAL SURFACES

The next level of software has only recently been built into hardware. This level refers to surfaces represented by polygons. *Tiling*, or polygon filling, is commonly available in hardware and software. Hidden surface removal is included. Sometimes anti-aliasing of polygon edges is provided to remove the distracting stair-steps, or jaggies, that result. Light sources can be incorporated into the rendered image, but

usually are point sources at infinity emitting white light.

PATCHES

The next level of sophistication represents surfaces as curved surface pieces called patches. This is still largely a software domain although hardware can be expected to appear soon. The most advanced software packages handle a variety of patch types. They also provide very sophisticated lighting models with multiple-colored lights, and distributed or point sources located either at infinity or in the scene being rendered.

Optical effects, such as transparency, translucency, refraction and reflection, are handled. Anti-aliasing is assumed. There is no practical software limit on scene complexity (such as the number of allowable polygons) built into software, but computation of highly complex scenes can take anywhere from 0.5 to 1.5 hours per frame on a supercomputer.

Research software provides even more features that produce greater realism in computer-generated imagery, such as articulated motion blur, depth of field, follow focus, constructive solid geometry and radiosity.

IMAGE PROCESSING

Image processing software has followed a separate path over the last 15 years or so. A variety of elaborate image processing software packages is now available. They provide functions such as convolution, Fourier transform, histogram, histogram equalization, edge detection, edge enhancement, noise reduction, thresholding, segmentation, bicubic and biquadratic warping, and resampling.

Appendix B. Visualization Environments: Short-term Potential

Many of these functions have been hard-wired into special boards. Only recently have processors become sufficiently powerful to make software competitive with hardware while maintaining generality. Image computers are capable of running both computer graphics and image processing software packages.

ANIMATION

A class of software that is only beginning to be appreciated but which is mandatory for visualization is that which controls animation. Animation, in its broadest sense, means movement; it frequently connotes the complex motion of many objects, possibly articulated, moving simultaneously and interacting with one another.

Animation is desirable for the visualization of complex processes. Basic animation control routines should be part of any standard visualization tool kit.

THE GLUE

A class of software appreciated by visualization professionals but not necessarily by scientists is the *glue* used to combine images generated or analyzed by the packages described above. These are the tools a user must have for convenience: picture compositing, picture

saving/restoring, fonts and text, resizing, rotation, moving, copying, hand retouching (*painting*), color manipulation, etc. Together, these functions comprise a *visualization environment system*, which is to visualization what an operating system is to general computing.

WINDOW SYSTEMS

Windowing systems are commonplace in black-and-white bit graphics and are being extended to color graphics. Visualization software will have to incorporate and be consistent with windowing concepts.

VOLUME VISUALIZATION

Volume graphics and imaging software is still rudimentary, and not yet implemented in hardware. Algorithms for rendering lines and curves into volumes have just become available. Similar algorithms for rendering surfaces and volumes into volume memories have yet to be presented. Anti-aliasing of volumes is not done, and hidden volume removal is unknown. The compositing of volumes is yet to be fully addressed. 3D paint programs (*sculpting programs*) have yet to be written. General utilities for arbitrary rotation and size change of volumes do not exist. In other words, there is much research to be done in this field.

B.2. Visualization Equipment

Supercomputers and workstations

Visualization tasks can be distributed across machine types. The four software functions of defining, constructing, rendering and displaying models may be implemented on different computer hardware. This hardware may even reside in different physical locations; that is, some of the hardware may be remote from the supercomputer.

The supercomputer-oriented approach often has the first three software functions implemented on the supercomputer, and usually in batch mode. The supercomputer creates images that are later displayed on a CRT or recorded onto film or video.

The workstation-oriented view uses the supercomputer for number-crunching model calculations or analyses, and uses specialized graphics workstations to render and display the results. Model development may also be done effectively on the workstation, or may be split between the workstation and the supercomputer.

Workstations encourage interactive computing, improving the productivity of the scientist. For example, the user can cancel bad runs, or can change parameters to examine different regions of interest. If the rendering is done on the workstation, the user can select arbitrary perspectives, lighting models, etc. A user with this type of dynamic control will avoid wasting computing resources on batch jobs to render erroneous or uninteresting data.

Workstations can work in stand-alone mode, or can serve as front-ends to supercomputers. The bandwidth required to send model data is usually modest, and is significantly smaller than that needed to handle rendered images. Remote support of direct interaction with the calculations may also be possible for certain applications. Remote *interactive* support of images rendered on a supercomputer and displayed on a workstation is prohibitively expensive at present.

There are distinct advantages to distributing software visualization tasks between supercomputers and workstations:

- Workstations offer a cost/effective complement to supercomputers.
- Workstations aid in the development of simulation models.
- Workstations can be used to guide the computations performed on supercomputers.

Appendix B. Visualization Environments: Short-term Potential

Supercomputer/Workstation Tradeoffs		
	Supercomputer with Display Only	Supercomputer/ Workstation Partitioning
Convenience	Effective only on-site	Permits effective remote support
Programming	Single environment	Two (often different) environments
	Limited graphics support	Excellent graphics support
Direct interaction (real time)	Possible for local investigator	More economical for local investigator and workstation
		Supports interactive rendering of results
Bandwidth (from supercomputer)	Very high speeds	Medium to low speeds
Costs	Wastes supercomputer for specialized graphics calculations	Achieves lower total cost by employing specialized graphics hardware.

Visualization Hardware

INPUT DEVICES

- **Sensors, cameras, computers**

Current digital input devices include supercomputers, satellites, medical scanners, seismic recorders, digitizing cameras, etc. The rapidly increasing bandwidth of these devices is enabling current work in volume visualization.

We can expect continued improvement in the resolution and bandwidth of input devices. For example, supercomputers will get faster and satellite resolutions will get higher. Real-time video digitizers already exist. Monochrome digital digitizers with a resolution of 2048x2048 are becoming quite fast although they are not yet real-time. Print quality input scanners (of flat work) are still quite expensive, but the prices can be expected to fall as digital technology cheapens and competing scanning technologies mature. CCD array input scanners will improve in resolution, and will become serious candidates for input devices in high-resolution work.

- **Interactive input devices**

Interactive devices are continually improving. Common 2D interactive devices include knobs, switches, pedals, mice and tablets. Tablets are the most general, and also need the most improvement; they need higher resolution, higher speed and more degrees of freedom.

3D interactive devices have begun to appear. 6D interactive devices are also available, providing the usual 3D positional information plus three degrees of orientation (yaw, pitch and roll) information. These will improve to offer higher resolution, higher speed and lower cost.

OUTPUT DEVICES

- **Raster displays**

Raster displays of 2D frame buffers have improved steadily over the years to offer more colors, higher resolutions and less flicker. A typical color raster display today offers a 1024x1024 pixel display at 60 frames per second and 24 bits of color per pixel (16 megacolors).

HDTV and 4:2:2 developments will affect visualization. HDTV is high definition television, a proposed standard which will offer larger, brighter, sharper pictures than currently available in video. Video is also going towards an all-digital format, designated 4:2:2, to standardize digital interconnections of diverse video products.

We can predict that color raster displays will evolve toward 2048x2048 pixels in the next several years. The displays themselves already exist in limited quantities; it is the computational bandwidth required to feed them that is still lacking. Black-and-white 2D raster displays already have resolutions greater than 2048x2048 with enough bandwidth to feed them. These displays will certainly reach even higher resolutions in the next five years.

Appendix B. Visualization Environments: Short-term Potential

- **Stereo displays**

Stereo displays are beginning to appear commercially. We can confidently predict that these will improve in screen size, resolution, brightness and availability. These displays will be quite helpful in volume visualization.

As with all stereo displays, a human viewer has to alter his visual system in some way to see the stereo. For example, a person might have to put on a pair of polarized glasses. If it is easy and economical to do this, and if consideration is given to people who already wear glasses, then individuals might find it sufficiently simple to let stereo become widespread.

- **Recorders**

Output devices are similarly improving. HDTV will spur the development of compatible recorders. Film recorders will become cheaper as the technology becomes cheaper and the competition matures. Should stereo become a widely accepted mode of presentation of volume visualizations, then stereo film and video standards will have to be developed.

WORKSTATIONS

- **Vector machines**

Fast vector machines are now commonplace and have extensive use in such areas as CAD/CAM and real-time 3D design. Recently they have been improved to offer color vectors and perfect end-matching. Frame buffers have been added so surface raster graphics may be combined with vector displays.

- **Surface machines**

Fast surface machines are just about to arrive. They exist in simplified forms already; they exist in more advanced states as firmware in special machines. Chips are now being built to speed up certain aspects of surface rendering, particularly the tiling of polygons. By 1990, full hardware support of surface graphics will be available, offering rendering features such as texture mapping, bump mapping, anti-aliasing, reflections, transparency and shadows.

Vector machines will initially serve as powerful real-time front ends to surface machines. Eventually, surface machines will be cheap enough and fast enough to permit scientists to do real-time design using surfaces rather than lines.

- **Image processing machines**

Fast planar machines have existed for some time; these machines contain special boards for certain aspects of image processing, such as fast Fourier transforms. Faster versions are now becoming available that have wider processing capabilities and higher resolutions. In fact, the notion of a general purpose image processor is becoming common; it is one which can implement any image processing algorithm as a program.

Observations of the way scientists are using visualization today suggest that a three-tiered model environment is evolving. Each model is distinguished by the costs and capabilities of its visualization technologies.

Visualization in Scientific Computing

B.3. Visualization Facilities

Three-tiered facilities model

Observations of the way scientists are using visualization today suggest that a three-tiered model environment is evolving. Each model is distinguished by the costs and capabilities of its visualization technologies.

It is natural to evaluate each model solely on its hardware's cost and power to compute. Four other attributes, however, equally affect scientific productivity: bandwidth, location, software support and administration structure. These attributes are defined for each of the three models in our environment.

Resource distribution provides better user control

Our model environment is predicated on the assumption that a scientist wants as direct a visual connection to his computation as possible. While supercomputers (model A) provide scientists with a powerful number-crunching tool for generating data, they do not currently do graphics; they fill arrays with information that somehow get piped to a display device. Workstations (models B and C), however, give scientists more control over their visual output. A workstation typically addresses its display memory the same way it addresses regular memory, incurring essentially no hardware overhead to display computed results.

Scientists should be able to select either the more expensive workstations with powerful visualization potential (model B) or the less expensive ones (model C), while maintaining network connections to the larger machines (model A) to do computations when necessary. This interdependency can work quite well; for example, a scientist can calculate 20-60 frames of a simulation sequence on a supercomputer, download the images to a workstation to create a mini-movie, and then play back the sequence at any speed under local control.

To make this work effectively, we need to improve the transfer of data to and from the main computation device. Networking then becomes a more critical factor than computer power in helping scientists do better science.

Appendix B. Visualization Environments: Short-term Potential

Visualization Facility Three-Tiered Hierarchy			
Attributes	Models		
	Model A	Model B	Model C
Hardware	Supercomputer or Super image computer	Mini-supercomputer or Image computer	Advanced workstations (Mini/micro image computers)
Bandwidth (potential interactive rates, bits/second)	$>10^9$	$10^7 - 10^8$	$10^3 - 10^6$
Location (where users interact with the display screen)	Machine room (at the Center)	Laboratory on a high-speed LAN (local area network)	Laboratory on a national/regional network (wide area network)
Software (in addition to discipline-specific data generation and processing)	<ul style="list-style-type: none"> • Commercial packages for output only (no steering) • Research required to develop interactive steering capabilities 	<ul style="list-style-type: none"> • Commercial packages are mostly output only • Some interaction is becoming available • Research required to improve discipline-specific interaction 	<ul style="list-style-type: none"> • Commercial packages and tools are widely available for both computation and interaction • Research required in languages, operating systems and networking
Administration	<ul style="list-style-type: none"> • Strength: Support staff • Weakness: Centralization 	<ul style="list-style-type: none"> • Strength: Discipline-specific visualization goals • Weakness: Small support staff 	<ul style="list-style-type: none"> • Strength: Decentralization • Weakness: No support staff

Assumptions

1. *In Model A, supercomputers and super image computers have equivalent power. The latter machine, although not commercially available today except in the form of a special-purpose flight simulator, will provide the specialized processing necessary for real-time volume visualization.*
2. *Model B assumes that mini-supercomputers and image computers have equivalent power.*
3. *Model C assumes advanced workstations are equivalent to mini/micro image computers in power.*
4. *Additional models D, E and F, which correspond to personal computers, alphanumeric CRT terminals and batch output, respectively, also exist. They do not represent advanced visualization technology, so they are not included in our model environment. Note however that type F has been used to produce a great deal of animation for both the scientific and commercial entertainment industries for the past 20 years.*

Only 10 percent of the supercomputer centers' budgets would be required to set up multiple instances of model B and C environments in-house — to become computer super centers — to accommodate a variety of user needs.

Visualization in Scientific Computing

Computer super centers

The total cost of setting up a model A facility is 100 times that of model B; model B costs approximately 10 times that of model C. Since 10 percent of a supercomputer center's budget (model A) can buy a large number of advanced workstations and a mini-supercomputer or two, it becomes clear that model A subsumes model B which subsumes model C.

Only 10 percent of the supercomputer centers' budgets would be required to set up multiple instances of model B and C environments in-house — to become *computer super centers* — to accommodate a variety of user needs. If supercomputer centers buy in quantity, then not only do hardware costs go down, but space in which the equipment resides is less expensive per square foot, visualization software is more affordable and networking becomes routine.

Accessibility to standard hardware, portable software and functioning networks would also enable scientists to share their work with colleagues at similarly equipped remote sites.

Three-Tiered Facilities Model: A Transportation System Analogy

Supercomputer centers are analogous to airline companies. Most people cannot justify personal jet aircraft, so they learn to deal with an administrative and temporal structure in order to satisfy their transportation needs. Likewise, scientists cannot afford supercomputers and specialized visualization tools, so they accept the fact that these resources must be governed in a centralized environment.

Conventional computer centers can be compared with public transportation in an increasingly car-centered culture. People

having special travel requirements buy their own cars; scientists wanting visualization have had to set up their own discipline-specific small computer centers.

The proposed three-tiered model environment advocates centralization of facilities only for the cases where there is no good decentralized solution. Decentralization requires improvements in documentation, computer-aided instruction, utility development and portable consultants.

Roadblocks to Establishing a Visualization Environment

Visualization is poorly supported by research funding, creating a number of short-term problems that must be solved before an ideal environment can exist.

AT SUPERCOMPUTER CENTERS

1. Supercomputing technologies are grossly out-of-balance with existing visualization technologies. Many more numbers are generated than can be understood.
2. It is harder to use supercomputers for visualization than small computers.
 - a. Visualization software is primitive, non-existent or very expensive.
 - b. Transfer of images between supercomputers and workstations is difficult.
 - c. File transfer protocols over networks have not been optimized to handle large binary files.
 - d. There is not yet a second tier of support processors for interactive visualization (such as advanced workstations or image computers) to fill the major gap that exists between workstations and supercomputers.
 - e. Hard copy equipment (film and video recorders, optical disks) is inadequate.
 - f. The typical supercomputer operating system is user hostile; that is, it is difficult to edit files, debug programs, etc.
3. Centers have few graphics experts on staff.
4. Centers have inadequate training labs for visualization.
5. Media production equipment, if any, is in a different building with a different director, and with no administrative or electronic connection.

AT CONVENTIONAL COMPUTER CENTERS

1. Centers have few graphics experts on staff.
2. Centers have inadequate training labs for

visualization.

3. Centers may/may not have commercially available visualization software, as it is very expensive.
4. Mainframes have primitive connections, if any, to workstations; conventional connections have too low a bandwidth for graphics.
5. Workstations use media (disks, tapes) which are unreadable by minis and mainframes, making it impossible to share information locally where networks do not exist.
6. Centers are reluctant to attach needed off-brand equipment, even if possible, for maintenance reasons.
7. Mainframe operating systems are not optimized for graphics, and are often crashed by incompletely tested and debugged drivers.
8. Media production equipment, if any, is in a different building with a different director, and with no administrative or electronic connection.

AT LABORATORY COMPUTING FACILITIES

A laboratory computing facility is the domain of an individual scientist or a small group of collaborating scientists. The scientist is responsible for doing both science and visualization. These scientists do visualization on their own, and just about everything else.

1. They must seek funding for visualization development, which is currently very difficult.
2. They must set up and maintain an independent computer center.
3. They must write their own visualization software and device drivers.
4. They must also find time to do science.
5. To publish their results in an accredited academic journal, they must spend time reducing their visualization methodology and insights to numbers and symbols.

Visualization in Scientific Computing

B.4. Network Usage

Temporary methods Today's networks are too slow for use in visualization. There are some techniques in place, however, to assist today's scientist. Provided the user is located near computational facilities, demands for high bandwidth can be reduced by a combination of techniques that include off-peak transmission of images, compression of images, reconstruction of images from abstract representations, and local generation of images.

Visualization and Networks

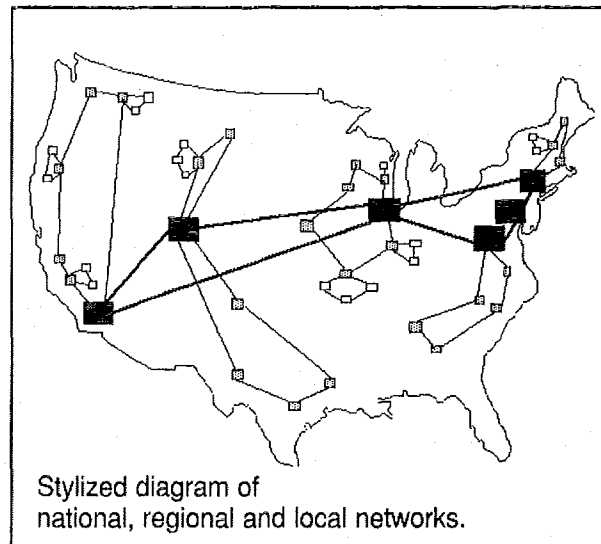
NSFNet

NSFNet, a three-tiered network communications system for the support of scientific computing, is evolving under the sponsorship of the NSF. Much work remains to be done to facilitate the process of linking scientists and engineers to appropriate computing resources, but the establishment of NSFNet is a major step toward this goal.

1. National networks link supercomputer centers.
2. Regional networks connect geographically-related institutions to nearby supercomputer sites.
3. Campus networks connect scientists to the regional networks and, ultimately, to the national supercomputer network.

Next generation wide-band networks

Wider band networks are needed for effective scientific communication involving images. The next generation of wide-band networks is probably three to four years away. Fiber digital



data interconnect (FDDI) chip sets capable of 100 megabit speeds are being designed but will probably not be available in quantity until 1988 or 1989. The establishment of a science network utilizing this technology would require an additional 2 to 3 more years for funding and completion.

Appendix C. Visualization Environments: Long-term Goals

Appendix C

Visualization Environments: Long-term Goals

The United States today faces a formidable challenge — to build the distribution networks, workstations and human interfaces necessary to communicate scientific and engineering knowledge. The magnitude of this task is difficult to see because our scientific and cultural biases in favor of traditional research methodologies would have us think this network is currently in place. Or, we assume that improved networking will soon alleviate the problems of transporting scientific simulations and engineering designs to the scientific/engineering marketplace. However, the problem is no more solved today than the muddy, impassable roads of France in the 1880's constituted an adequate transportation system. Modern industrialized France still needed to build government-mandated networks of new roads and railroads in the years prior to World War I.

In the 1980's, visualization communication in the United States is hobbled by lack of standards, mired in the intellectual frustration of making interconnections across incompatible media, and held up at the gateways by disparate transmission protocols never designed with visualization in mind. Visual communication cannot be shared among users across a distributed network of incompatible workstations with idiosyncratic user interfaces and no common layering of portable software or hardware.

Scientific communication is changing. Traditionally scientific publication has meant English language print, line drawings and a few static images. Increasingly, science cannot be done in print; in fact, essentially all of the visualization research opportunities described in this report require visualization networking and visualization-compatible electronic media for publication.

This appendix identifies the needs for visualization-enhanced scientific communication in a number of areas:

- Visual publication
- Human/computer interfaces
- Televisualization
- Visualization software and hardware
- Standards
- Education

Contemporary scientific communications media are predominantly language-oriented. Printed media are coupled weakly, if at all, to the visual world of space-time. By contrast, half the human neocortex is devoted to visual information processing. In other words, current scientific communication leaves out half—the right half—of the brain.

Visualization in Scientific Computing

C.1. Scientific Visual Communication and Publication

The importance of visual communication

Edison's invention of the electric light bulb was predicated upon the construction of a distribution system for electricity. Edison said he "wanted to make electricity so cheap that only the rich could afford candles." Today, visualization can be generated at isolated workstations, but methods for distributing the visual information that results have yet to be designed, prototyped, produced and publicized.

Contemporary scientific communications media are predominantly language-oriented. Printed media are coupled weakly, if at all, to the visual world of space-time. By contrast, half the human neocortex is devoted to visual information processing. In other words, current scientific communication leaves out half — the right half — of the brain. An integral part of our visualization task is to facilitate visual communication from scientist to scientist, from engineer to engineer, through the intermediary of visualization-compatible communications media.

While *interaction* today describes the scientist's potential to direct his computation and synthesize new insights dynamically, *interaction* has a social meaning as well. "Do you see what I see?" one researcher asks another. In this way, hypotheses can be tested and validated or falsified in minutes instead of years. Changing the scale and pace of visualization alone would affect research profoundly. But we can predict with certainty that such changes in modality will also lead to immense conceptual advances as well.

Visual publications

Scientific research requiring computationally intensive visualization is in danger of becoming Babelized and thereby incommunicable. Much of modern scientific research cannot be expressed in print — DNA sequences, molecular models, medical imaging scans, brain maps, simulated flights through a terrain, simulations of fluid flow, and so on. If poorly communicated, such research cannot stimulate new discoveries and new disciplines.

The end results of selected visualization — photographs, films and videotapes — are what most people see. With the exception of flight simulator trainees and video game players, all visualization seen by those not involved in producing it is one-way; it is non-interactive. A scientist cannot publish "interaction" in a journal.

*Visualization and science go hand in hand as partners.
No one ever expected Gutenberg to be Shakespeare
as well.*

Appendix C. Visualization Environments: Long-term Goals

Electronic media, such as videotapes, laser disks, optical disks and floppy disks, are now necessary for the publication and dissemination of mathematical models, processing algorithms, computer programs, experimental data and scientific simulations. The reviewer and the reader will need to test models, evaluate algorithms and execute programs themselves, interactively, without an author's assistance. Scientific publication needs to be extended to make use of visualization-compatible media.

Visualization language: a new cultural development

Reading and writing were only democratized in the past 100 years, and are the accepted communication tools for scientists and engineers today. A new communication tool, visualization, in time will also be democratized and embraced by the great researchers of the future.

Visualization: The Infant Giant	
Communications Media	Number of Years Old
Sight	5×10^8
Speech	5×10^5
Writing	5×10^3
Print broadcasting	5×10^2
Visual broadcasting	5×10^1
Visualization	5×10^0

The introduction of visualization technology will profoundly transform the way science is communicated and will facilitate the commission of large-scale engineering projects. Visualization and science go hand in hand as partners. No one ever expected Gutenberg to be Shakespeare as well. Perhaps we will not have to wait 150 years this time for the geniuses to catch up to the technology.

C.2. Human/Computer Interface

Human interaction styles Among the modes of human/computer interaction seldom available in large computing environments, and therefore unfamiliar to many scientists, is the style popular in artificial intelligence computing environments. The activities and features that characterize this style are:

- Rapid prototyping
- Graphical monitoring
- Layered abstractions

These approaches contrast strongly with the traditional means of interfacing with computer systems that are used by many physical scientists:

- Inflexible languages
- Batch compilation
- Repeated linking of the entire program structure
- Batch-mode running of linked programs with no interactive monitoring
- Restriction to primitive numerical data structures

Human/computer interaction is still in its infancy. The techniques discussed in this section are just the beginning of what could be accomplished, even within the context of existing capabilities.

Rapid prototyping

Small pieces of code are changed repeatedly and executed using an interpreted language invoked from an environment such as a text editor. To accomplish this, one needs a language that can be incrementally compiled and an editor with powerful system interfaces. Since linking and loading of the program are not required, one typically achieves modification and debugging cycles on the order of a few seconds.

When testing different approaches to human/computer interaction, numerous experiments with short turn-around times like this are appropriate to explore the possibilities. Inflexible environments discourage interesting exploration.

Appendix C. Visualization Environments: Long-term Goals

Graphical monitoring

High resolution, multi-window, multi-process displays can be easily used to monitor the progress of a program by showing pictorial data in various windows. This technique permits the scientist to monitor simultaneous events in real time, find incorrect or misunderstood behaviors, and correct them.

Layered abstractions

The use of languages that accept user-supplied macros and permit the hiding of data and program details are essential aspects of modern programming methods. These capabilities support the extension of a language to handle new application requirements without changing the basic environment. With graphics and images, such languages tend to be *object-oriented*, with self-describing data types.

Scientific understanding

The field of visualization requires a new formulation of the problems of transforming symbolic and numerical descriptions into graphic renderings and simulated behaviors. What if the symbols and simulated behaviors are in the minds of scientists? How can the computer get high-level access to the mental imagery of what the researcher wants to express? New concepts are required that define human capabilities for analogical and metaphorical visualization, transformations of visually representable concepts, and generation of intuitive understandings that lead to conceptual scientific breakthroughs.

Concurrent processing of symbolic information and visual imagery

Many modes of thought require concurrent processing of symbolic information and visual imagery. For example, when we read a textual description, perhaps of a room in a mountain inn, we can often visualize the scene. Textbooks provide both text and supporting illustrations, enabling us to reference either or both for better comprehension.

Knowledge-based systems, however, do not currently have this capability. The concurrent processing of verbal and visual information will be of critical importance if computer-based intelligent systems are to capture any significant portion of the world's knowledge as archived in published literature.

Visualization in Scientific Computing

Contributions need to be made to the development and growth of computer-based, scientific communication utilities. Among contemporary systems, there are two restrictions that we feel can be partially alleviated:

- The absence of visual modes of thought; specifically, the inability to co-reference between textual descriptions and associated illustrations in published literature.
- The inability to integrate scientific knowledge from multiple sources.

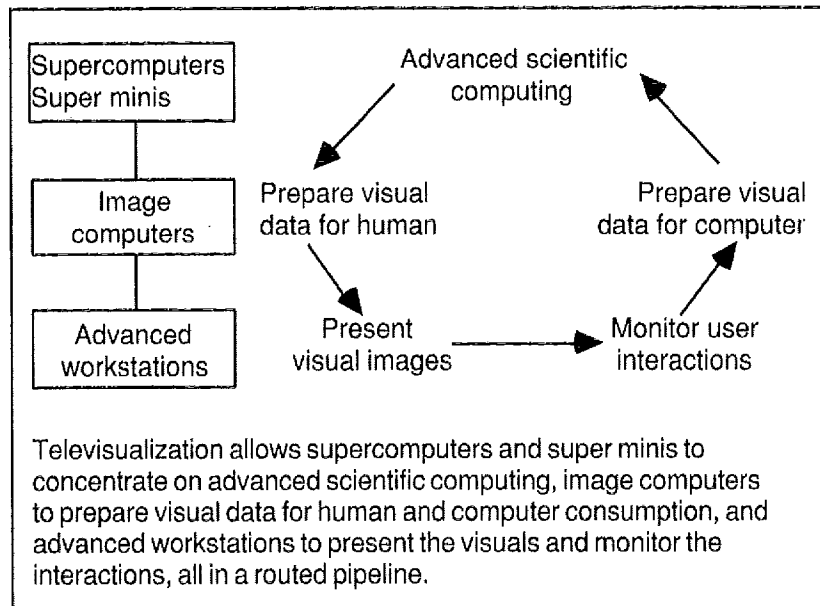
The application of networks to visualization is termed *televisualization*. The authors encourage federally funded network researchers to acknowledge the needs of televisualization. A high-level team approach, emanating from the current network initiative, is required, including experts in networking and visualization hardware/software.

Appendix C. Visualization Environments: Long-term Goals

C.3. Televisualization

What is televisualization? The sheer scale of graphics and image data sets challenges the current bandwidth and interactivity of networks. Networks that handle screenfuls of textual information exist and work well; network nodes are simply gateways that neither add nor detract from the quality of the message. But, a 512x512 pixel image (x8 bits/pixel) has approximately 100 times more information than a screen of text (25 rows x 80 characters/row).

The application of networks to visualization is termed *televisualization*. The authors encourage federally funded network researchers to acknowledge the needs of televisualization. A high-level team approach, emanating from the current network initiative, is required, including experts in networking and visualization hardware and software. Televisualization encompasses much more than the transfer of text (for example, electronic mail) and gateway protocol decoding. It involves the transfer of images, which entails compression, decompression, rendering, recognizing and interpreting.



Visualization in Scientific Computing

Televisualization requires a major enhancement over existing network capabilities in the following areas:

- Increased data rates
- Value-added processing at nodes
- Interaction capabilities

Data rates

In the future, computing will be highly distributed. Not only will computation be distributed across multiple processors in parallel architectures, but it will be distributed across multiple cooperating users via networks. To perform the visualization integral to scientific computation routinely, we need to focus attention on improving network bandwidth. Visualization files are large, so passing them efficiently from station to station or supercomputer to supercomputer will require wide-band communication channels.

Bandwidth much higher than currently available will be required to ship images among users. Gigabit speeds are sufficient to pass volumes of the current size of 256x256x256 pixels (at 4 bytes per pixel, 2^{26} bytes or 64 megabytes per volume). Within several years, this rate will have to be extended to 1 gigabyte per second channels to handle larger volume sizes, with image size of 4 gigabytes or more.

Value-added processing at nodes

Computers process text and numbers in main memory, occasionally transmitting some of them to peripherals. Images, however, must often be transferred to special memories for rendering and/or viewing. Each instance of transferring and processing an image aims to increase its visualization value to the scientist. This ability to process images at various nodes along a network embraces the central concept of distributed computing.

Transmitting 3D images or volumes over an wide-area network is not practical; transmission costs and times are too high. Instead, we need a scheme whereby we can balance transmission costs with local computing costs. It makes more sense to send model data over networks and then render or reconstruct it at the other end. This, however, presumes that there is appropriate equipment at the other end and that the software, although running on a variety of equipment types, is compatible. Hence, televisualization supports reduced transmission costs by advocating the availability of more local visualization hardware.

Appendix C. Visualization Environments: Long-term Goals

Interaction

A televisualization network of *image passing* between machines is analogous to the software paradigm of *message passing* between layers of processes running on a machine. This type of network cannot be achieved using conventional FORTRAN subroutine calls. Significant software development and protocol standardization are necessary to bring televisualization to the discipline sciences.

Visualization access and sharing

A number of user issues involving networks are not yet well understood and need to be studied in order for scientists to take advantage of appropriate computing resources, whether local or remote.

1. How does software prototyped on one scientist's workstation migrate over a network to more powerful computing facilities?
2. How can we use networks to share software among geographically dispersed scientists within a specialized area of investigation?
3. What are the security constraints that must be imposed so that scientists may feel free to send their data over a network to a remote site?
4. How can the considerations of cost and convenience be balanced to satisfy the computing needs of the scientific enterprise?
5. What qualitatively different science requires interactivity? What science can be done at 50 megabytes per second that cannot be done at 50 megabits per second?

A major problem to be solved as we enter the 1990's is the need for portable, robust, useful visualization software.

Visualization in Scientific Computing

C.4. Visualization Software and Hardware

LAYERED VISUALIZATION SOFTWARE

A major problem to be solved as we enter the 1990's is the need for portable, robust, useful visualization software. Currently there is no standard set of visualization tools and graphics interface conventions appropriate for complex tasks, the definition of which is an open research problem.

Based on existing technologies and trends, there is a need for a family of compatible visualization software systems of varying levels of functionality and performance. There are many benefits to be gained from this layered software approach.

- The software can be easily adapted to the variety of data sources, computers and display devices.
- The software can be tailored to the needs of diverse applications; not all applications require the highest performance visualization hardware, software and communication technologies.
- The software allows more techniques, algorithms, etc. to be added as they become available.
- Different software layers can be distributed to different machines in a distributed network; that is, layers can be distributed over one or several machines so that varying capabilities and their associated costs can be tailored to individual situations.
- The software is more economically feasible to develop for a limited scientific market as a common set of tools than as turnkey systems tailored to specific visualization applications.
- The software can accommodate the differing budgets, equipment preferences and requirements of a broad constituency of users.
- The demands of the scientific computing community can be coherently addressed in unison.

A layered software approach requires definition of the functional layers, message passing protocols and communication bandwidth/volume information. Layers would communicate by message passing rather than the traditional procedural control passing schemes. A layer would import messages from its neighboring layers and export messages to its neighboring layers; a null or empty implementation of a layer would merely pass messages between its neighbors. There are four functional layers in the model we are proposing:

1. Database and data source management
2. Modeling
3. Rendering
4. Display

Appendix C. Visualization Environments: Long-term Goals

Database and data source management

Data generated by scientific application programs (data sources) may be stored in commercial database management systems (DBMS), in special-purpose systems, or not at all. Implementers of large scientific applications are not known for the attention they pay to managing large volumes of output. Regardless, this data, which is the output of scientific computations, becomes the input to visualization programs and must therefore be managed more rigorously.

Modeling layer

The modeling layer provides software for both classification and surface modeling.

- *Classification* software either detects boundaries or classifies voxels. Boundary detection of 2D images detects and models boundaries between regions of relatively uniform properties. Voxel classification of 3D images categorizes each voxel by its composition; for example, in a volume reconstruction of a CT scan, voxels would be classified as fat, bone or muscle. Boundary detection within volumes is a special case of classification which determines whether a voxel is inside or outside an area.
- *Modeling* applies a geometric description of an object and its attributes to data extracted from the database. Principal geometric modeling schema include enumerated or cellular, faceted, functional or sweep, and procedural or constructive solid geometry. The choice of a model is affected by the fundamental nature of the data, the style and quality of rendering desired, and the costs of data storage and communication.

Rendering layer

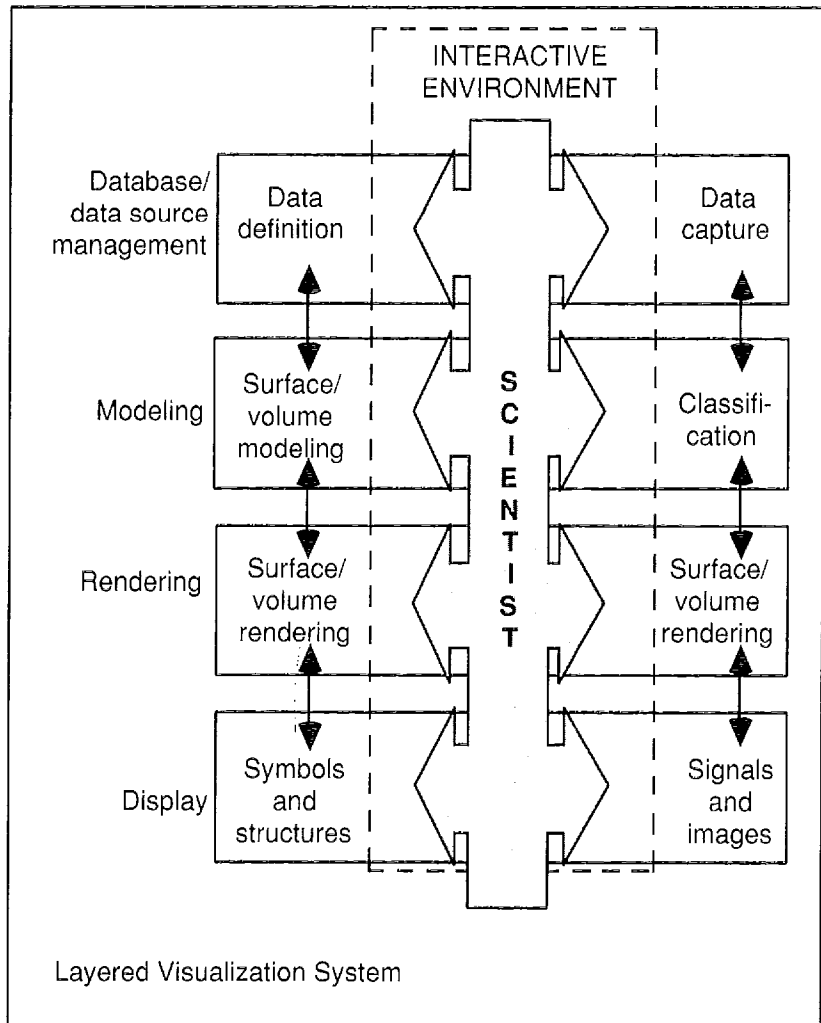
The rendering layer renders either surfaces or volumes:

- *Surface rendering* software accepts 3D geometric descriptions with corresponding surface attributes and combines this information with viewing parameters (lighting, orientation, viewing angle, etc.) to produce a metafile for export to the display layer.
- *Volume rendering* software accepts 3D volume descriptions with surface and volume attributes and combines this information with viewing parameters to produce a metafile.

The rendering level must be able to accept a variety of surface and volume formats, provide rendering at various qualities and costs, utilize techniques such as incremental rendering for improved interaction and throughput, and provide a *background mode* for such computationally intensive techniques as ray-tracing and radiosity.

Visualization in Scientific Computing

Display layer The display layer contains device drivers that deal with the idiosyncrasies of I/O devices to communicate user interaction information or computer-generated images, to and from other software layers. More specifically, the drivers are capable of sending user input to different layers in the system, or receiving metafiles of 2D and 3D image descriptions from the other layers for display on a monitor (vector, raster and stereo) or for recording to a hard copy device (paper, transparency, slide, film/video and stereo film/video).



Appendix C. Visualization Environments: Long-term Goals

Interaction handling

Interactivity provides a natural means for a scientist to communicate with data by manipulating its visual representation. The most exciting potential of visualization tools is not the still color photographs or movies a scientist can produce to portray the results of his simulation or computation, but the ability to test and develop code interactively. The scientist should be able to spot visual anomalies while computing and immediately *steer*, or modify the calculations, to test out new theories. Interactivity puts the scientist back in the computing loop; the scientist gets to *drive* the scientific discovery process.

Distributed processing environment

The ability to port these layers to different machine types would allow a user to choose how hardware costs, communication costs and non-local computing costs should be allocated, an approach which has a long history of success in graphics. For example, a minimal user workstation might be no more than an adequate display (8 bits deep, 24 bit look-up table, RGB monitor) on an inexpensive personal computer; only the display layer would be resident on the station. In more powerful systems, layers associated with rendering, modeling, database and data source management might reside on the workstation. Those layers not in the workstation would exist elsewhere, such as on a mainframe or supercomputer.

Interactivity issues, especially in a distributed environment, deserve some separate discussion. Messages associated with interaction must be accepted by the display layer from the user and passed to the appropriate layer for action. For example, a user request for contrast adjustment would be handled at the display layer level; a user request for rotation of the displayed object would be passed to the rendering layer.

Long-term Visualization Software Problems

We can identify a number of long-term technical problems that impact all types of visualization computing environments:

1. *Neither commonly available languages nor operating systems exist for visualization.*
All large operating systems are optimized for string and numerical processing, with the exception of those used in television studios. The UNIX operating system, for example, is based on a teletype mode of interaction. The only operating systems that integrate visualization well are available exclusively on personal computers and workstations.
2. *Existing software is not easily vectorized or ported to parallel computers.*
Image algorithms make massive use of lists, trees and shared memory (the image), which frustrates vectorization and parallelization. New algorithms must be developed.
3. *Documentation of visualization software on paper leads to 500-page door-stop manuals.*
The documentation needs to be *interactive* and *visual*.
4. *Visualization standards have not been supported by the U.S. Government.*
The only standard in force is a foreign one, based on 15-year old workstation capabilities, which is incomplete with respect to 3D and raster graphics, and which has no discernible conceptual model to extend to multi-dimensional, time-varying imagery. While the rapid development of visualization hardware makes standards for software premature, the authors believe that a focal point about which standards can develop is required.
5. *Image/message protocols do not exist across machine types.*
Although many workstations handle image data well internally and among machines of the same type on networks, there is no image/message protocol that works across a wide range of computers.

Note that we differentiate between the central computer with the majority of the computational power and the specialized visualization computer, just as human anatomy has differentiated between the specialized visual cortex and the brain.

Appendix C. Visualization Environments: Long-term Goals

SPECIALIZED VISUALIZATION HARDWARE

Getting data to and from the main computation device is a complex operation; graphics devices need to send and receive newly computed data in real time. Transforming raw data from its initially computed form to an effective image often requires expensive techniques such as hidden surface removal, visibility determination, ray-tracing, and so on.

Much needs to be done in the design of interfaces and special graphics hardware to make all these operations possible on the same time scale, presumably fractions of a second, to match the speed at which scientists can assimilate graphic information from a visualization device.

General computer technology

A number of technologies promise faster computational speeds for computers. These include faster silicon from submicron chip-making, gallium arsenide, optical computers, neural net architectures and other highly parallel computers.

Perhaps the most exciting recent advance is the achievement of superconductivity at more normal temperatures. The cheap cooling this promises would make extremely *dense* computers possible; the term *dense* means *faster and cheaper*. This development will make the fabrication of 3D chips thermodynamically feasible, chips integrated in 3 dimensions as compared to today's 2.

Image computers and super image computers

Image computers for surface and volume visualization have only recently appeared on the market. These machines have very large memories that make 256x256x256 volumes easy to manipulate and have sufficiently powerful processors to compute raster volumes of this size quickly.

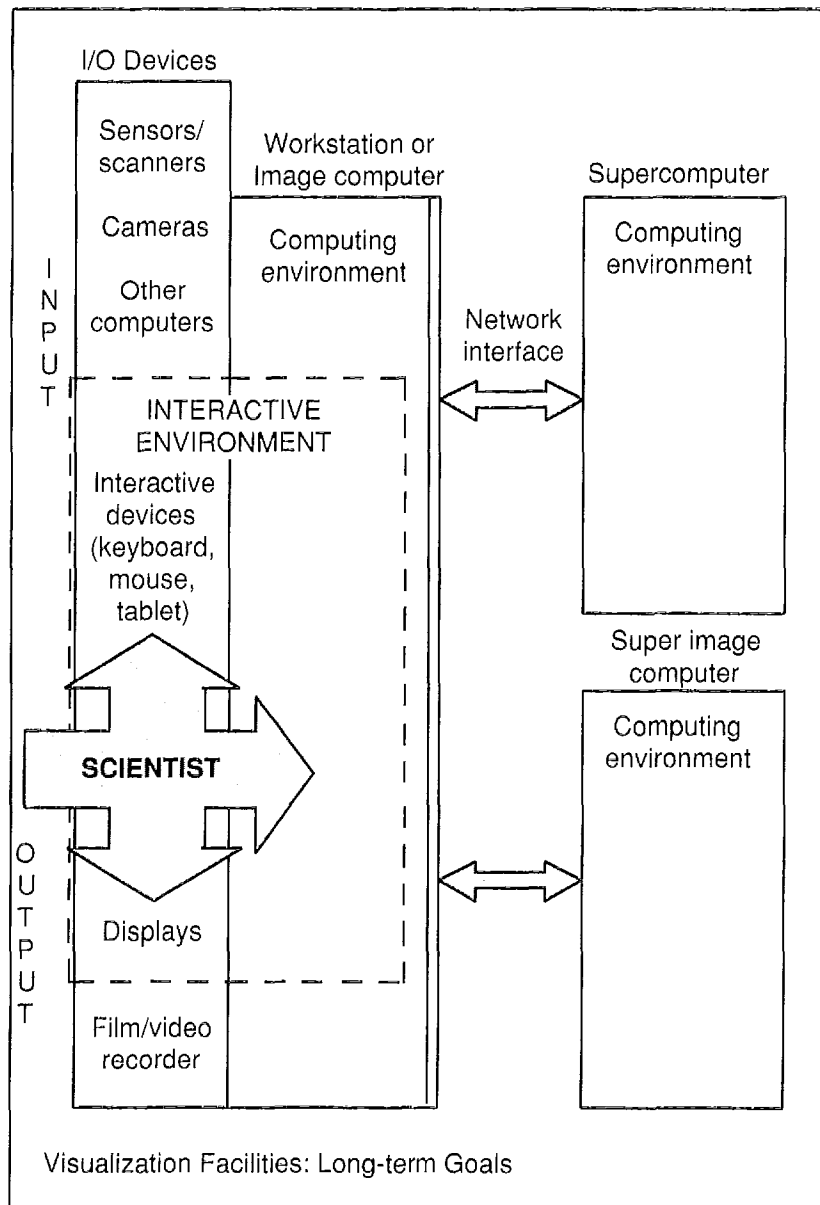
Over the next 2-5 years, highly parallel architectures will evolve to the point where image computers will be able to approach speeds offering real-time interaction. There will also be an increase in the volumes which can be handled, up to such convenient sizes as 1024x1024x1024 voxels; at 4 bytes/voxel, that is 2^{32} , or 4 gigabytes, of memory.

The speed of the image computer refers to the rate at which computations can fill and manipulate raster memory; computations that generate the data that gets displayed would run on another processor, such as a supercomputer. Note that we differentiate between the central computer with the majority of the computational power and the specialized visualization computer, just as human anatomy has differentiated between the specialized visual cortex and the brain.

Visualization in Scientific Computing

Mass storage

Mass storage media are evolving to handle the vast amounts of data required by volume visualization. Gigabyte storage devices are readily available at megabyte per second access rates. Volume visualization will require terabyte storage devices at gigabyte per second access rates. In the meantime, increased data storage capacity will become available with optical techniques, and increased access speeds with parallel drives and parallel heads.



In truth, standards are absolutely essential to the widespread adoption of a technology, particularly any technology which is a communications medium.

Appendix C. Visualization Environments: Long-term Goals

C.5. Standards

Attitudes towards standards

It is a curious thing that telephones, telegraphs, television, power transmission and radio are all nationally standardized; yet, standards for the most powerful and pervasive communications link to come are not yet a national concern.

Some people consider standards the busy work of small minds who come along after the leaders and crusaders have shown the way. Others view them as shoddy restrictions, created *de facto* by some large commercial concern with less than grand intentions. In truth, standards are absolutely essential to the widespread adoption of a technology, particularly any technology which is a communications medium.

NTSC is the American national television standard. It was designed before there was color television, so color was crammed into the standard. The international move to establish a new standard (HDTV) attempts to cross national lines to define a new standard based on a much more advanced notion of the television medium.

PostScript is a recent candidate for a robust computer-displayed text standard. It is well thought out, elegant and extensible as our understanding of computer fonts and text improves; specifically, it allows the addition of color and 3D fonts.

The need for good visualization standards

We believe that it is too early in the life of the visualization discipline to impose standards, but we can recommend categories of visualization which should be standardized in order for tool makers and tool users to move forward on a national scale. By developing a structure within which standards can be established, we can discourage further fragmentation — to prevent the Balkanization of visualization.

Visualization and Standards

There are several graphics *standards* in place. Several of them are being enhanced to include newer aspects of our understanding. All of them were born in the early, vector-based days of computer graphics when all transformations were those provided by 4x4 matrices and when surface graphics was too expensive to even consider.

Most research centers have disregarded these standards as being too restrictive or old-fashioned to consider. Consequently, there currently exist an abundance of graphics packages which are only standard where they were written. Some of them differ in quite simple ways: Is the coordinate system in which objects are defined left-handed or right-handed? Does the y-axis on a display screen point up or down? Does the z-axis point in or out? Are the three orientation degrees of freedom called yaw, pitch and roll, or are they called azimuth, elevation and twist (or every possible combination thereof)? Are matrices multiplied from the left or the right? Is the aspect ratio defined in terms of height to width, or width to height? Is the viewing angle the half angle or the whole angle? These differences make the interchange of models between research centers very difficult.

There is little commonality about the representation of splines, patches and polygons. There is less ability to handle modern graphics modeling notions, such as procedural and stochastic models. There is no strong capability for surface rendering parameters in any

standard, although some standards are currently being extended to handle some of these: colors, lights, shadows, textures, bumps, environments, material types, matte versus glossy finish, etc. Facilities for animation are largely missing from all current standards. Transformations different from the 4x4 matrix ones are generally not handled.

There are simply no standards yet for volume visualization as this field is too new. Any efforts here would have to be easily extensible as understanding of the field evolves.

A very common thing to do in graphics is to ship picture files to other facilities. There is, however, no standard for a picture file. Very large picture files and volume files over insufficient network channels have forced the development of picture compression/decompression schemes. There are many of these and, of course, no standardization.

High-bandwidth networks will require new, visualization-capable, networking protocols. These still have not been proposed. Rather, each industry is beginning to think about defining industry-wide standards. For example, Picture Archiving and Communication Systems (PACS) protocols in medicine are being discussed. Similarly, CAD/CAM has generated several interchange protocols. Remote sensing has developed still others, and so has the print industry. Little emphasis has been placed on inter-industry compatibility.

Appendix C. Visualization Environments: Long-term Goals

The software, hardware, and televisualization network models discussed previously serve as a foundation from which methodologies can develop. Technological evolution will select from these methodologies; some will die out while others will be adapted as *de facto* standards. These various models, however, must be nurtured to give some direction and definition to the methodologies that will evolve. Specifically, we propose that attention be directed towards nurturing the following ideas:

1. Scientific research projects, in addition to their other merits, should document their visualization methodologies and distribute them to the visualization community at large.
2. Supercomputer centers should adopt at the earliest practical time a visualization standard for use with their equipment, which could then serve as a model for national standardization.
3. The scientific community should raise the level of standards consideration across competitive lines and even national lines and encourage a long-term view of visualization.

We must raise the importance of a good visualization standard (or a set of standards) to such a high level of national priority that the best tool users and builders in the country will want to take part.

Computer scientists must have access to visualization technology, and computational scientists must learn to think visually.

Visualization in Scientific Computing

C.6. Education

Visual education

A major educational reform must take place on two fronts. Visualization technologies must exist and be used by tool makers and tool users. Computer scientists must have access to visualization technology, and computational scientists must learn to think visually. At present, there are roadblocks:

- The Association for Computing Machinery's (ACM) approved computer science curriculum lists computer graphics as merely one of many optional topics; image processing is not mentioned at all.
- Engineering accreditors do not require computer graphics or image processing.
- Many engineering school deans are unaware of the importance of visualization and/or cannot justify the hardware and software expenses involved.
- There are approximately the same number of tenured faculty teaching computer graphics in American universities today as there were 15 years ago, and they are roughly the same people.
- Scientists, while educated to read and write, are not taught to produce or communicate with visuals.

Visual training, if any, is left to the commercial television and entertainment industries. Our culture is simply reluctant to accept, much less favor, any level of non-verbal communication in science. Our culture talks of the insight necessary for science, but only favors that form of *seeing* which can be reduced to typed symbols and statistics.

Academic and scientific reward systems

Today's academic and scientific reward systems overwhelmingly favor those who manipulate numbers and symbols. We have used computers to replace much phenomenological observation and rely mainly on symbol processing to gain insight. We trust numbers and symbols; they are easy to compare and publish. Academia has an enormous installed base of numerical data acquisition and analysis machinery, word processing and printing technology — all in support of getting science in print.

Publishing and grants, and therefore tenure, have rarely come to those whose productivity depends upon or produces visualization results. Scholarly work is evaluated by superiors by counting the number of journal articles published; publications are text, and visual media do not count. Funding itself is based on the careful preparation and evaluation of proposals, which are documents full of words and numbers.

Industry relies on the research community, especially the university-based research community, for advances in basic technology, new methodologies, and the next generation of trained personnel. Researchers depend upon industry for the instrumentation needed for their investigations and experiments.

Appendix D. Industrial Competitiveness

Appendix D

Industrial Competitiveness

A strategic economic objective of the United States today is to maintain its leadership in technology. For scientists and engineers to gain fundamental insights, significant investments are being made in advanced scientific computing. The authors of this report believe that support for visualization is the most effective way to leverage this investment in national competitiveness.

Competitiveness requires a cycle of development that runs faster, and is better targeted, than that of one's competitors. The cycle itself depends on a push-pull interaction between industrialists and researchers. Industry relies on the research community, especially the university-based research community, for advances in basic technology, new methodologies, and the next generation of trained personnel. Researchers depend upon industry for the instrumentation needed for their investigations and experiments.

If researchers cannot gain access to state-of-the-art instrumentation, research deteriorates. If they are unable to provide new insights and solutions, they cannot pull industry with new products and industrial processes. The development of research tools more advanced than those commonly available in the commercial marketplace — whether developed in academia or industry — has to be supported, or the productive cycle is stymied, industrial advancement slowed, and the competitive edge lost.

Researchers in the discipline sciences and related engineering fields are now starting to utilize many of the visualization techniques developed over the last several decades. They are, however, inhibited by their lack of experience in techniques for computer manipulation of the visual concepts that are so fundamental to their disciplines. More than ever, they now require a means to interact with designs and simulate new concepts before building the products that keep industry competitive.

This proposed ViSC initiative focuses on two goals: (1) producing a generation of scientists and engineers trained in visualization techniques who can migrate into industrial and educational environments to keep them competitive, and (2) developing a set of visualization tools, programs and techniques that most fully leverage the research investment over all the disciplines.

D.1. Visualization in the Design and Engineering Process

Trial-and-error manual methods of scientific and engineering design

The trial-and-error process is essential to scientific and engineering progress. Models of the behavior of air flow over a new aircraft wing, of the structure of a new chemical, or of the structural and aeronautical integrity of a new spacecraft design, must each be tried and tested before being used. In the past, this has largely been a manual process. A physical model was (and in some instances still is) constructed; for example, a small model of a new aircraft is built and subjected to wind tunnel tests to verify its characteristics. In the automotive industry, automotive crash tests are used to determine the safety of new cars.

These manual methods, by their very nature, are extremely costly, time-consuming and inadequate, since many tests are required to ensure reliability and consistency. This trial-and-error process can be shortened by computer modeling and simulation, making the product available more rapidly, with greater reliability, and with greater benefit to industry and our country.

Computer-assisted scientific and engineering design

The modern, computer-assisted design process has three major steps.

- *Visual design and definition:* Scientists and/or engineers specify a visual definition of an object and its properties. That is, they construct a spatial or geometric model of the object of design; for example, they specify the shape of an aircraft wing, the structure of a molecule, or the design of a mechanical part.
- *Modeling/analysis:* Scientists and/or engineers conduct detailed analyses, often using a supercomputer, of the fundamental physical laws and differential equations governing their designs. A numerical study models the behavior of a particular design in a simulated environment, such as the air flow over a wing in a simulated tunnel, the minimum energy of a molecular configuration, or a structural load on a new automobile part.
- *Visual interpretation:* Scientists and/or engineers visually examine the results of a simulation combined with the original model, such as particle flows over an airplane wing, energy contours of a new molecule, or color stress contours on a mechanical part. Combining the original design with the results of analytical simulation yields insights that are not possible with conventional, manual techniques. Properties that are not visible

Visualization is not limited to supercomputers and high-performance multi-processor machines. Design and discovery processes can also be carried out on smaller problems using mini-supercomputers and advanced workstations, but visualization is most blatantly absent at the high end.

Appendix D. Industrial Competitiveness

or are difficult to measure, such as temperature and stress distribution on a mechanical part or turbulence over a new airfoil, become evident with the use of various computer display techniques.

These fundamental processes are a unique outgrowth of the computer and computer graphics industries. Many application areas are on the threshold of a major transformation from the older, more costly and more time-consuming methods of actual model construction and testing to newer computer simulation techniques. The ViSC initiative should ensure that scientific and engineering research centers make this inevitable transformation rapidly.

Visualization facilities for scientific and engineering design

The national supercomputer centers' program provides funding for the development of the second step of the three-step design process defined above: design, modeling/analysis and interpretation. The ViSC initiative addresses the first and last steps.

DARPA, the Defense Advanced Research Projects Agency, under its strategic computing initiative, is supporting the development or upgrade of several high-performance machines designed for image processing and computer vision, such as the Butterfly Parallel Processor, the WARP Systolic Array Processor and the Connection Machine. Again, funding is provided for the second step of this three-step design process. Little is being done to give these machines high-speed display capabilities. Visualization hardware and software would benefit both users, whose goal is to process and display images, and researchers, who are developing image processing and understanding algorithms.

Visualization is not limited to supercomputers and high-performance multi-processor machines. Design and discovery processes can also be carried out on smaller problems using mini-supercomputers and advanced workstations, but visualization is most blatantly absent at the high end. At all levels, a visualization initiative will help the country's scientists and engineers gain the fundamental insights needed for technological leadership.

The Japanese computer graphics industry is coming up fast, attracted primarily by a rapidly growing world market and a better awareness of the strategic significance of developmental tools using computer graphics.

Visualization in Scientific Computing

D.2. International Competitiveness of the US's Visualization Industry

Japanese computer graphics market research

International competitiveness refers primarily to America's technical and industrial rivalry with Japan. In the 1985 *Report on Japanese Computer Graphics Industry and Market* by two leading computer graphics consultants Carl Machover and Laurin Herr, it was shown that as late as 1983 imported products accounted for more than 50 percent of the approximately \$730M domestic Japanese market for computer graphics. In the same year, Japanese manufacturers were estimated to have exported more than \$125M worth of graphic display terminals, plotters, monitors and flat displays. By 1988, imported products are expected to drop to less than 40 percent of the estimated \$4.1B Japanese domestic market. In the same period, Japanese exports are projected to increase almost 900 percent.

The bottom-line assessment of this comprehensive report is that America continues to lead the computer graphics world in technology and therefore in market share. This includes a very strong presence within Japan, the world's second largest market for these products. However, the larger Japanese firms are increasingly turning their attention towards computer graphics-related technologies. As a result, the Japanese computer graphics industry is coming up fast, attracted primarily by a rapidly growing world market and a better awareness of the strategic significance of developmental tools using computer graphics.

The market forecasts in this report are based on data collected in Japan during 1983-1985. While reliable, they are somewhat dated. There is always a ± 15 percent margin for error in any study about as young and rapidly evolving a subject as Japanese computer graphics. Moreover, unforeseen developments, such as the more than 40-percent drop in the value of the dollar during the past two years, have further skewed the forecasts made in 1985. This major currency shift has made Japanese exports less price-competitive, American exports more price-competitive, and direct investment in American corporations an unexpectedly cost-effective way for Japanese firms to gain rapid access to American technologies and markets. However, the basic assumptions underlying the numeric projections published in 1985 still hold true.

Japanese companies are expected to continue their domination of the global market for CRT monitors and emerging flat display technologies. Their enormous strength in the consumer television business provides a valuable economic base for continued development of high-resolution color display devices at competitive prices.

The emergence of Japan as the home of leading-edge manufacturing and design logically implies a future market demand within Japan for leading-edge computer graphics applications systems.

Appendix D. Industrial Competitiveness

International growth of the Japanese computer graphics market

Attempts by Japanese companies to increase their penetration of the world computer graphics market will initially focus on peripheral devices, since peripherals involve relatively little software development, usually require only printed circuit board level integration, and can be manufactured using highly automated electro-mechanical assembly processes. Extrapolating from the rush by many Japanese manufacturers into computer graphics in the 1983-1985 period, the character of Japanese entry behavior in other markets, and the relatively small size of the domestic Japanese market, we can expect the rapid development of excess Japanese manufacturing capacity in peripheral devices.

This will lead to aggressive export efforts and steep declines in the per-unit cost of similar devices around the world. This, in turn, will squeeze the profits of existing vendors and will cause a shakeout of many of the small and medium-sized computer graphics firms in the United States.

The 1983-1985 Japanese exports of graphics terminals failed to gain wide acceptance in America. This was due to the rapidly shifting nature of the market and the existence of many small American vendors with excellent technology. However, the increasing stabilization of the international computer graphics market around various hardware and software standards is expected to stimulate the emergence of various commodity-like hardware products, a development thought to favor Japanese manufacturing prowess and institutional stamina in any future struggle for international market share.

The future of Japanese market demand

The emergence of Japan as the home of leading-edge manufacturing and design logically implies a future market demand within Japan for leading-edge computer graphics applications systems. These systems might, in some cases, be more sophisticated than those developed by American vendors to satisfy American users.

Such market demand will increasingly be met by domestic Japanese suppliers. One reason is the improved technical capabilities of Japanese firms. Another reason is the American vendors' dependence on Japanese distributors, which characteristically limits feedback from Japanese users. If American vendors don't get closer to their Japanese customers — and if America does not generally increase its access to Japanese technical and scientific information — American superiority in the crucial systems and software areas are expected to seriously erode during the next decade. Also relevant in this context are the surge of Japanese investments in industrial R&D,

The United States must advocate a forward-looking defense of their market position in the face of strenuous Japanese efforts to improve their capabilities in this and related fields.

Visualization in Scientific Computing

efforts by the Japanese government and industry to promote better software development, and the well-honed Japanese talent for adaptive improvement of imported concepts.

America: Complacency or forward-looking defense

There are, and will continue to be, global market leadership advantages accruing to American vendors. However, complacency is dangerous. The United States must advocate a *forward-looking defense* of their market position in the face of strenuous Japanese efforts to improve their capabilities in this and related fields.

Already Japanese firms dominate the global markets for video equipment and 35mm photo cameras, which are the heart of general-purpose image recording and reproduction today. Japanese firms also have a very strong base in photographic films, videotape, optical storage and image transmission. Japanese color printing is widely considered to be among the best in the world. Almost every leading Japanese firm in the video, photography, film and printing industries is involved in some aspect of computer graphics development, either at the commercial or at the R&D levels.

For several years, the American computer graphics industry has been seriously dependent upon Japanese suppliers for color monitors, flat displays, CCD imaging devices, thermal printing devices, and semiconductor memories. Japanese firms also have very considerable strengths in related areas such as machine tools and robotics, not to mention their growing expertise in supercomputing, artificial intelligence and microprocessors.

Sophisticated Japanese VLSI graphic controllers have been widely adopted by American graphics manufacturers in the recent past. While these pioneering Japanese chips have lost the competitive edge to a new generation of more powerful, programmable chips introduced by American firms within the past year, the Japanese can be expected to counter with new products of their own in this competitive market segment.

A significant percentage of the peripheral devices used in large, American-made, computer graphics systems are manufactured by American firms. Many of the smaller American-made systems, however, are increasingly relying on Japanese peripherals, such as small scanners, desktop laser printers, small pen plotters, ink jet printers, small tablets, mice, keyboards, magnetic disks and optical disks. While profitability in larger peripherals is still quite good, the market for smaller peripherals is growing much faster in terms of unit sales and total value.

Appendix D. Industrial Competitiveness

Japanese computer manufacturers have quietly penetrated the low end of the American market in computing platforms for graphics systems with IBM/PC and PC/AT clones, sold under their own or other's brand names. Within Japan, American manufacturers of personal computers hold only negligible market shares. Most Japanese graphics-based personal computers are built around the 16-bit NEC-9800 series.

American engineering workstation vendors are still the recognized global market leaders at this time. However, several Japanese companies have also developed their own 32-bit UNIX workstations. In fact, NEC introduced its first workstation into the United States' market at the end of 1986, reportedly prepared to sell "industry-standard products in large volume, with good quality, and at a very aggressive price," according to NEC executives. NEC can plan such a price-oriented market entry strategy this late in the game because it is so much larger and more vertically integrated than top American workstation vendors.

Japanese government and quasi-government sponsored research

Clearly, one of the reasons for rapid Japanese progress now is the strong underpinning provided by government-sponsored research in the past, especially the crucial national projects in VLSI and Pattern Information Processing Systems (PIPS). Research results from quasi-governmental labs have also been widely disseminated throughout the large Japanese electronics firms. These include the work by NHK (the Japanese national broadcasting company) in high-definition television or HDTV, and the work by NTT (national telephone company) in facsimile. Several current government-sponsored research efforts in Japan, such as the SIGMA software engineering project, the Supercomputing Project, and the Fifth Generation Project, may also yield significant benefits for the Japanese computer graphics industry in the next decade.

Another important reason for rapid Japanese advances in computer graphics is superior Japanese efforts to gather and analyze technical and commercial information from around the world. Most of this industrial intelligence work is performed by Japanese industry, but the government, too, has always taken a leadership role in this area.

For example, in 1985 MITI (the Japanese Ministry of International Trade and Industry) organized a committee of experts to comprehensively assess the Japanese market and industrial base in computer graphics. In 1986 MITI successfully persuaded the rival NICOGRAPH Association and the Japan Computer Graphics Association to merge "for the sake of Japan." According to participants, MITI argued that a single industry association speaking

Computer graphics is one of the few outstanding areas of clear American technical and market superiority over Pacific Rim rivals.

Visualization in Scientific Computing

with one voice would be better able to gather and disseminate information, coordinate interaction with foreign groups, and foster the growth of the domestic industry.

The newly formed Nippon Computer Graphics Association and MITI are now jointly raising funds to create a comprehensive video and text curriculum about computer graphics for use in professional schools throughout Japan. This curriculum is expected to include considerable amounts of materials solicited, purchased or produced in America. It is clearly aimed at acquainting a new generation of Japanese professionals with the potential of computer graphics, thereby stimulating the further growth and development of the Japanese computer graphics industry.

Significant benefits for America's international competitiveness

Computer graphics is one of the few outstanding areas of clear American technical and market superiority over Pacific Rim rivals. It is one of the few such areas still remaining, and one that might well be dominated in the near future by all the new technologies and new applications coming out of the Japanese R&D establishment.

Loss of this superiority would seriously weaken America's general R&D capabilities with unacceptable consequences for many sectors of American industry. Sustained efforts to promote further progress in this field are urgently needed to compete with the fairly consistent Japanese strategy of technical self-improvement. In this context, a national ViSC initiative represents an historic opportunity to build on America's strengths. Also, it promises significant benefits for America's international competitiveness far into the 21st century.

Appendix E

ViSC Videotape: List of Contributors

The *Visualization in Scientific Computing* report consists of this printed document and a two-hour videotape that illustrates pioneering efforts in visualization today. It is very frustrating to try to verbally describe the impact of visual information, and the contents of these tapes serve to illustrate and enhance these printed words.

The tapes were produced, in part, with a grant from the ACM SIGGRAPH organization. They are being distributed as issues 28 and 29 of *SIGGRAPH Video Review*. Ordering information appears on page ii of this report.

The editors received many excellent submissions from which only two hours could be selected for inclusion in this report. Additional materials may appear in subsequent issues of the *SIGGRAPH Video Review*. The authors encourage computational scientists and engineers who rely on visualization tools for their work to show their films and tapes at conferences, to submit them to SIGGRAPH for publication, and to encourage their own discipline science organizations to recognize film and video as viable forms of publication.

Credits: These tapes were edited by Thomas A. DeFanti (University of Illinois at Chicago), Maxine D. Brown (University of Illinois at Chicago) and Bruce H. McCormick (Texas A&M University). Coordination and title slide preparation by Bill Gale (University of Illinois at Chicago). Production by Raul Zaritsky (E&C Media, Chicago, Illinois).

Visualization in Scientific Computing

E.1. SIGGRAPH Video Review Issue 28: Visualization Domain


1. L.A. — THE MOVIE



Contact: Kevin Hussey or Jeff Hall, Jet Propulsion Laboratory (JPL), 4800 Oak Grove Dr., M/S 168-522, Pasadena, CA 91109. **Credits:** Kevin Hussey, Bob Mortensen,

Jeff Hall. **Technical Notes:** Computational science/engineering discipline: Image Processing. Description of problem: Three-dimensional perspectives of data (Landsat, TM, DMA topographic data) with motion to create a simulated fly-through. Hardware: VAX 8600 for computation; MicroVAX-controlled video recording equipment mastered on Betacam. Software: Software was developed by the authors at JPL. Time: Computation time was 130 CPU hours on the VAX 8600. © Copyright 1986, Jet Propulsion Laboratory/NASA.

2. INSTABILITIES IN SUPERSONIC FLOWS




Contact: Michael L. Norman, National Center for Supercomputing Applications (NCSA), 605 E. Springfield Ave., Champaign, IL 61820. **Credits:** Michael L.

Norman, Philip E. Hardee, David A. Clarke, Donna J. Cox. **Technical Notes:** Computational science/engineering discipline: Astrophysics/Fluid Dynamics. Description of problem: Development of a kink instability in a supersonic gas flow which is initially collimated and laminar. Hardware: Cray X-MP/48, Silicon Graphics IRIS, Abekas. Software: ZEUS astrophysics code performed the simulation and rasterization of images. ICARE software was used to modify the color palette. Time: 10 CPU hours were used to compute the fluid simulation as 200,000 zones for

10,000 time steps. © Copyright 1987, National Center for Supercomputing Applications, Board of Trustees, University of Illinois at Urbana-Champaign.

3. CALTECH STUDIES IN MODELING AND MOTION




Contact: Alan H. Barr, California Institute of Technology (CalTech), Computer Science Department 256-80, Pasadena, CA 91125.

Credits: Al Barr, Ronen

Barzel, Tim Kay, Jed Lengyel, John Platt, John Snyder, Brian Von Herzen. **Technical Notes:** Computational science/engineering discipline: Physically-Based Modeling. Description of problem: Modeling based on physical simulations using *dynamic constraints*: (1) Self-assembly and motion of chains, (2) Self-assembly of space structures, (3) Elastic material, and (4) Molecular modeling. Hardware: Motion was computed on a Symbolics 3675 Lisp machine. Images were rendered on HP 9000 Series 300 computers. Software: Software was developed by the authors at CalTech. © Copyright 1987, California Institute of Technology.

4. EVOLUTION OF STRUCTURE IN THE UNIVERSE



Contact: Joan M. Centrella, Department of Physics and Atmospheric Science, Drexel University, Philadelphia, PA 19104. **Credits:** Joan M. Centrella (Drexel)

and Craig Upson (formerly at Digital Productions and now at NCSA, Champaign, IL). **Technical Notes:** Computational science/engineering discipline: Astrophysics. Description of problem:

Appendix E. ViSC Videotape: List of Contributors

The evolution of large-scale structure in the universe. The model produces long, dense filaments and nearly empty voids which are similar to features observed in the distribution of galaxies. Hardware: Cray X-MP/22, film recorders, E&S and IMI workstations. Software: Digital Productions DP3D proprietary software. Time: Approximately 50 Cray X-MP hours. © Copyright 1985, Joan M. Centrella.

5. DYNAMIC CRACK PROPAGATION WITH STEP-FUNCTION STRESS LOADING




Contact: Robert B. Haber, National Center for Supercomputing Applications (NCSA), 605 E. Springfield Ave., Champaign, IL 61820.

Credits: Hyun M. Koh,

Hae-Sung Lee and Robert B. Haber (Principal Investigator). **Technical Notes:** Computational science/engineering discipline: Solid Mechanics. Description of problem: Finite element simulation of dynamic fracture and elastic wave propagation in a brittle rectangular plate subjected to sudden stress loading. Hardware: A Cray X-MP/48 was used for the finite element simulations. A Raster Technologies One/380 attached to a VAXstation II was used for graphics post-processing. An Abekas videostore was used for video animation recording. Software: Custom finite element software. Raster Tech ONELIB host subroutine software. Time: The finite element analysis took approximately 10 Cray hours total. Graphics post-production took approximately 3 days on the RT380/Abekas. © Copyright 1987, National Center for Supercomputing Applications, Board of Trustees, University of Illinois at Urbana-Champaign.

6. NUMERICAL SIMULATION OF A THUNDERSTORM OUTFLOW



Contact: Robert Wilhelmson, National Center for Supercomputing Applications (NCSA), 605 E. Springfield Ave., Champaign, IL 61820.

Credits: Kelvin Droegemeier and Robert Wilhelmson. **Technical Notes:** Computational science/engineering discipline: Atmospheric Science. Description of problem: Simulation of the outflow which is similar to a cold front that splits between precipitating thunderstorms. Hardware: Cray X-MP/22. Software: Digital Productions DP3D proprietary software and 2D numerical cloud model written by Kelvin Droegemeier. Time: 40 Cray-1 equivalent hours.

7. SCIENTIFIC VISUALIZATION — SCIENCE DATA SYSTEMS GROUP/JPL



Contact: Robert S. Wolff or Betsy Asher, Jet Propulsion Laboratory (JPL), 4800 Oak Grove Dr., M/S 168-522, Pasadena, CA 91109.


Credits: Robert S. Wolff,

Kevin Hussey, Betsy Asher, Dennis Blakey, Peter Farson, Jeff Hall, Hemanshu Lakhani, Linda Roy.

Technical Notes: Computational science/engineering discipline: Global cloud dynamics. Description of problem: A demonstration of atmospheric data visualization in 2 and 3 dimensions. Software: 3D perspective projection software developed as part of JPL/VICAR at MIPL by Bob Mortensen and Kevin Hussey. Cubic interpolation software developed as part of JPL/VICAR at DIAL by Jeff Hall. Image processing and animation performed at MIPL and DIAL using JPL/VICAR software by Jeff Hall and Kevin Hussey. © Copyright 1987, Jet Propulsion Laboratory/NASA.

Visualization in Scientific Computing


8. POLIOVIRUS



Contact: Arthur J. Olson, Department of Molecular Biology MB5, Research Institute of Scripps Clinic, La Jolla, CA 92037. **Credits:** Arthur J. Olson and Dan Bloch.

Technical Notes: Computational science/engineering discipline: Structural Molecular Biology. Description of problem: The atomic resolution structure of the intact poliovirus capsid is visualized. The film shows the viral architecture and subunit structure. Hardware: VAX 11/750 and Evans & Sutherland Multi-Picture System. Images were filmed directly off the screen with a Mitchell 16mm animation camera. Software: GRAMPS, a graphics language interpreter by T.J. O'Donnell and A.J. Olson, and GRANNY, a molecular modeling package by M.L. Connolly and A.J. Olson. © Copyright 1987, Research Institute of Scripps Clinic.

9. INERTIAL CONFINEMENT FUSION




Contact: Nelson Max, Lawrence Livermore National Laboratory, P.O. Box 808, M/S L-301, Livermore, CA 94550.

Credits: Computer animation by Nelson Max,

optical printing by John Blunden. **Technical Notes:** Description of problem: Simulation of target implosion, ignition and burn during the laser-driven inertial confinement fusion of deuterium-tritium fuel. Hardware: CDC 7600, Sperry Univac V75, Dicomed D48. Time: The vast majority of time was spent on the Sperry Univac V75 minicomputer determining the shading and driving the Dicomed D48 color film recorder, and on creating optical effects, like the expanding plasma, from computer graphics images.

10. RPI SCIENTIFIC VISUALIZATION




Contact: Mary N. Johnson, Assistant Director, RPI Center for Interactive Computer Graphics (CICG), CII-7105, Troy, NY 12180.

Credits: Phil Getto,

David Breen, Randy Bradley. **Technical Notes:** Description of problem: A*R*M (Another Robot Movie) simulates an idealized robot arm picking up bolts from a conveyor belt and placing them on an assembly block. This is one application of The Clockworks, a general-purpose visualization software tool capable of simulating any 3D structure that can be geometrically and mathematically modeled. Hardware: Data General DG Mv-10000, Raster Technologies One/380, Evans & Sutherland PS/300, Dunn 635 camera. Software: The Clockworks object-oriented image synthesis/simulation system developed by RPI/CICG. Time: Cray-1 hour equivalent ratio approximately 1:80+/-20. 1200 CPU hours were used for the A*R*M animation sequence, rendering time for other images ranged from 0:23 to 14:00. © Copyright 1985, RPI/CICG.

11. RIGID BODY DYNAMICS SIMULATIONS



Contact: James Hahn or Charles Csur, Ohio State University ACCAD, 1501 Neil Ave., Columbus, OH 43201. **Credits:**


Simulation software and animation by James Hahn.

Technical Notes: Computational science/engineering discipline: Dynamics (Mechanical Engineering), Computer Simulation. Description of problem: The dynamics of rigid bodies, including articulated bodies, are simulated. Interactions among many bodies are modeled (both instantaneous and continuous contact). Hardware: Symbolics 3600/3620 for animation calculation, Convex C1 for raster version.

Appendix E. ViSC Videotape: List of Contributors

Software: Rigid body dynamics simulation software by James Hahn, scan-line algorithm by Scott Dyer and Scott Whitman. Time: The simulation was approximately 0.2 Cray-1-hour equivalents. The raster version was approximately 11 Cray-1-hour equivalents. © Copyright 1987, James Hahn.

12. NASA/CFD HIGHLIGHTS




Contact: Thomas Lasinski, NASA Ames Research Center, M/S 258-5, Moffett Field, CA 94035. **Technical**

Notes: Computational science/engineering

discipline: Computational Fluid Dynamics and Computational Aerodynamics. Description of problem: This piece demonstrates an application of RIP (Remote Interactive Particle Tracer) to solve the partial differential equations (PDE's) governing fluid flow. The flow is calculated over the surface of a wing of an F16-like aircraft; the particle traces are colored using pressure as a scalar function. Also demonstrated are turbine hot streak analysis and 3D turbine geometry. Hardware: Cray X-MP, Silicon Graphics IRIS 2500T/3030, Lyon-Lamb animation system. Software: Specialized software to solve PDE's of fluid mechanics, and special graphics routines based on Silicon Graphics' software library.

13. COMPUTATIONAL FLUID DYNAMICS



Contact: Karl-Heinz A. Winkler, Los Alamos National Laboratory, P.O. Box 1663, M/S P-370, Los Alamos, NM 87545.


Technical Notes:

Computational

science/engineering discipline: Fluid Dynamics.

© Copyright 1986, Regents of the University of California.

14. AEROSPACE APPLICATIONS OF ADAM AND POSTPROCESSOR




Contact: Sharon L. Gallup, Mechanical Dynamics Inc., 3055 Plymouth Rd., Ann Arbor, MI 48105. **Credits:** Software development by Mechanical Dynamics Inc.

MBB Satellite Coupling problem developed by TEDAS GmbH, a distributor of Mechanical Dynamics. **Technical Notes:** Computational science/engineering discipline: Aerospace Engineering. Description of problem: This tape contains simulations of models for which engineering information, such as displacements, velocities, accelerations and reaction forces, have been calculated. Specifically, it contains a simulation of satellite docking mechanisms and maneuvers. Hardware: Software is supported on the MicroVAX II, VAX 11/process, various IBM computers including the IBM/RT, Cray, Apollo, Sun, Hewlett-Packard 9000/500, Silicon Graphics IRIS, Prime. Animation is provided on Apollo, Evans & Sutherland PS/300, Silicon Graphics, Sun-3, Tektronix 4115B and upward terminals. Software: Mechanical Dynamics' ADAMS and Postprocessor software. © Copyright 1987, ADAMS, Mechanical Dynamics Inc.

E.2. SIGGRAPH Video Review Issue 29: Visualization Systems

1. VOLUME VISUALIZATION WITH THE PIXAR IMAGE COMPUTER



Contact: Dana Batali,
Pixar, P.O. Box 13719,
San Rafael, CA 94913.
Credits: Weather data
courtesy of TASC. Gas
dynamics data courtesy of
Ronald K. Hanson,


Stanford High-Temperature Gas Dynamics Laboratory, Mechanical Engineering Department, Stanford University. Smoke data courtesy of NASA Ames, and Juan Agüí and Lambertus Hesselink of the Department of Aeronautics and Astronautics, Stanford University. Fluid dynamics data courtesy of Princeton University. Supersonic broomstick data courtesy of the University of Illinois Center for Supercomputer R&D. PATRAN stress analysis data courtesy of PDA Engineering. CT data of temporal bones and pelvis courtesy of Philips Medical Systems, Inc. CT data of chest and shoulders courtesy of Elliot Fishman, John Hopkins University. CT data of sea otter courtesy of Michael K. Stoskopf, Veterinary Medicine, John Hopkins University. **Technical Notes:** This demo tape contains several examples of the emerging computer graphics technique called Volume Visualization. Five principal volume visualization tools are used:

- *Rendering.* Conventional 3D rendering techniques render views for giving depth or motion information in conjunction with the filmloop tool below, or for generating fully-colored slices (cross sections) of 3D objects or data sets to use with the reconstruction (stacking) tool and cubetool below.
- *Reconstruction (Stacking).* 2D slices of a multi-dimensional object or data set are *stacked* into a 3D cube of data, which may or may not be fully occupied with the given data. The view point is arbitrary.
- *Cubetool.* The cubetool allows a user to

interactively move through a cube of data, generated with the stacking tool perhaps, using a mouse or a tablet. Real-time motion along each or all of the three principal axes is provided. The cube may be reoriented to have any face perpendicular to any axis.

- *Filmloop.* Views of a 3D volume are generated — generally not in real time — and stored in disjoint parts of a large memory. These are displayed in quick succession from the memory by remapping the video display to the disjoint parts. This remapping occurs at video rates (30 frames per second), giving the illusion of a film running in real time. Since the *film* is coming from memory, a user may interact with it in real time by single stepping, reversing the direction of motion, rocking, pausing indefinitely, or changing the speed of display. This tool is frequently used to edit a visualization exercise for subsequent videotaping.
 - *Classification.* Colors are assigned to volumes in a meaningful way.
- © Copyright 1987, Pixar.

2. CONNECTION MACHINE APPLICATIONS




Contact: Janet
MacLaren, Thinking
Machines Corporation,
245 First St., Cambridge,
MA 02142. **Technical
Notes:** This demo
explores various

Connection Machine applications, including analysis of aerial images and fluid dynamics. Hardware: Connection Machine CM-1.
© Copyright 1987, Thinking Machines Corporation.

Appendix E. ViSC Videotape: List of Contributors

3. IMAGE PROCESSING ON PIPE




Contact: Ernest W. Kent, Philips Laboratories, 345 Scarborough Rd., Briarcliff Manor, NY 10510.

Credits: Ernest W. Kent. (Work was done at the

National Bureau of Standards.) **Technical Notes:** Computational science/engineering discipline: Image Processing Architecture. Hardware: This demo illustrates the power and speed of the prototype PIPE (Pipeline Image Processing Engine) machine. Software: None.

4. PIXEL-PLANES 4 DEMONSTRATION TAPE




Contact: Henry Fuchs, University of North Carolina, Department of Computer Science, Sitterson Hall (083A), Chapel Hill, NC 27514.

Credits: Henry Fuchs,

John Poulton, John Austin, John Eyles, Trey Greer, Clare Durand, David Ellsworth, Mukesh Gupta, Steve Molnar, Brice Tebbs, Greg Turk. **Technical Notes:** Computational science/engineering discipline: Computer Graphics Hardware. Description of problem: Molecular modeling, 3D medical imaging, dynamic (real-time) shadow generation. Hardware: Fourth generation of the Pixel-planes experimental graphics system. The system's novel feature is a frame buffer of 512x512 pixels x 72 bits/pixel implemented on 2,048 custom logic-enhanced memory chips. The system renders about 35,000 full-screen, smooth-shaded, Z-buffered polygons per second; about 13,000 spheres per second. Software: The graphics software was developed by the University of North Carolina. It runs on a UNIX-based workstation which hosts the Pixel-plane's custom graphics hardware. © Copyright 1987, University of North Carolina at Chapel Hill.


5. REDISCOVER ENGINEERING



Contact: Gail Madison, Silicon Graphics Inc., 2011 Stierlin Rd., Mountain View, CA 94043. **Technical**

Notes: Description of problem: This demo reel contains a variety of Manufacturing Computer-Aided Engineering (MCAE) examples that run on the Silicon Graphics' IRIS superworkstations. © Copyright 1987, Silicon Graphics Inc.

6. SYNTHETIC HOLOGRAPHY



Contact: Timothy P. Browne, MIT Media Laboratory, E15-222, 20 Ames St., Cambridge, MA 02139. **Technical**

Notes: Description of problem: Under the direction of Steve Benton, the Spatial Imaging Group's aim is to develop and demonstrate the powerful potential of 3D imaging as an efficient means for communications. One of their recent achievements is the invention of a large-scale, multi-colored, computer-graphic hologram. © Copyright 1987, MIT Media Laboratory.

7. CARTOGRAPHIC MODELING SYSTEM




Contact: Andrew J. Hanson, SRI International, 333 Ravenswood Ave., M/S EK-239, Menlo Park, CA 94025. **Credits:** L. Quam, A.J. Hanson, A. Pentland, P. Fua. SRI

International Artificial Intelligence Center. **Technical Notes:** Computational science/engineering discipline: Interactive Image Understanding. Description of problem: Extract semantic information from digitized imagery. Hardware: Symbolics 3600 Lisp Machine with Hi-Res Frame Buffer. Software: Special-purpose Lisp

Visualization in Scientific Computing

systems. Time: Simulated scenes were generated in 2 minutes/frame on the Symbolics 3600 (equivalent to a VAX 11/780). © Copyright 1987, SRI International.

8. HP 9000 SRX




Contact: Glenda McCall, Hewlett-Packard, Fort Collins Systems Division, 3404 East Harmony Rd., Fort Collins, CO 80525.

Technical Notes: This tape describes the

functionality of the HP SRX computer graphics system. HP 9000 SRX footage, courtesy of Hewlett-Packard Company. © Copyright 1987, Hewlett-Packard Company.

9. VIDEO REPORT ON THE COMPUTER GRAPHICS INDUSTRY



Contact: Frost & Sullivan Inc., 106 Fulton St., Dept. R-1, New York, NY 10038. **Credits:**

Produced, written and directed by Raul Zaritsky of E&C Media, Chicago,

IL, and Laurin Herr of Pacific Interface, New York, NY. **Technical Notes:** Description of problem:

This is an excerpt from a series of one-hour video reports published by Frost & Sullivan and hosted by Laurin Herr. This series of reports survey the visualization hardware and software available on the market today. © Copyright 1987, Frost & Sullivan Inc.

Visualization in Scientific Computing at NSF

Visualization in Scientific Computing at NSF

Although it does not yet have a formal initiative for Visualization in Scientific Computing, the National Science Foundation (NSF) supports several programs that would be responsive to proposals with a visualization component. Direct inquiries to one of the following program directors by contacting:

National Science Foundation
1800 G Street, N.W.
Washington, D.C. 20550

Symbolic and Numeric Computation
Computational Mathematics
Design Theory and Methodology
Software Systems
Computational Science and Engineering/New Technologies
Applied Mathematics
Computational Engineering
Interactive Systems
Computer Integrated Engineering

Bob Caviness
Ray Chin
Susan Finger
Harry Hedges
Richard Hirsh
Bill Lakin
George Lea
Laurence Rosenberg
Tony Woo



Local Groups Currently Operating

Eugene Miya
Bay Area ACM/SIGGRAPH
NASA Ames Research Center
Mail Stop 233-14
Moffett Field, CA 94035
(415) 694-6453

Maria Mezzina
Chicago ACM/SIGGRAPH
School of Art Inst. of Chicago
Columbus Dr. & Jackson Blvd.
Chicago, IL 60603
(312) 243-9542

Dick Moberg Eric Podietz
Delaware Valley ACM/SIGGRAPH
P.O. Box 1954
Philadelphia, PA 19105
(215) 923-3299

Dave Miller
Denver/Boulder ACM/SIGGRAPH
830 Magnolia Street
Denver, CO 80220
(303) 388-3231

Guy Fontenier Bernard Dresner
France ACM/SIGGRAPH
17, rue de la Baume
75008 Paris, FRANCE
(331) 43.59.77.55

Harold Santo
Lisbon ACM/SIGGRAPH
Centro de Mecanica e Engenharia
Estruturais
da Universidade Technica de Lisboa CMEST
Instituto Superior Technico-1096 Lisboa Codex
PORTUGAL
80 15 79

Jim Banister
Los Angeles ACM/SIGGRAPH
TRW MS RI 1062
T Space Park
Redondo Beach, CA 90278
(213) 536-3605

Bob Hirshon
New England ACM/SIGGRAPH
P.O. Box 194
Bedford, MA 01730
(617) 738-0421

Carol Chiani Ken Glickfeld
New York ACM/SIGGRAPH
Video Computer Animation Workshop
503 Broadway
New York, NY 10012
(212) 925-9356

Jeremy Rowe
Phoenix ACM/SIGGRAPH
Media Systems, Box JAV
Tempe, AZ 85287
(602) 965-5046
(602) 839-3713

Douglas Dixon
Princeton ACM/SIGGRAPH
David Sarnoff Research Center
CN 5300
Princeton, NJ 08543-5300
(609) 734-3176

William Haake
Rochester ACM/SIGGRAPH
Taylor Hall
University of Rochester
Rochester, NY 14627
(716) 275-9173

Mike Amron
San Diego ACM/SIGGRAPH
3311 Spa Street
San Diego, CA 92105
(619) 265-2092

Norm Cohen
Syracuse ACM/SIGGRAPH
P.O. Box 6154
Syracuse, NY 13217
(315) 471-0800

Garri Garripoli
Washington, D.C. ACM/SIGGRAPH
Circuit Studios
2121 Newport Place, N.W.
Washington, D.C. 20037
(301) 659-6275

Future Conference Dates

SIGGRAPH '88

August 1-5, 1988
Atlanta, Georgia
Andy Goodrich
(313) 763-4888
Adele Newton
(519) 888-4534

SIGGRAPH '89

July 31-August 4, 1989
Boston, Massachusetts
Branko Gerovac
(617) 480-6647

Christopher Herot
(617) 494-1400

SIGGRAPH '90

August 6-10, 1990
Dallas, Texas

SIGGRAPH '91

August 5-9, 1991
San Francisco, California

SIGGRAPH '93

August 2-6, 1993
Anaheim, California

