Why is Real-Time Volume Rendering No Longer a Year Away?

Organizer
Arie Kaufman, State University of New York at Stony Brook

Panelists
Marty Brady, Intel Corporation
Bill Lorensen, General Electric Corporation
Fred Kitson, Hewlett-Packard Labs
Hanspeter Pfister, MERL — A Mitsubishi Electric Research Lab

INTRODUCTION

Something new is happening in computer graphics. Volume rendering matured as a method for visualizing 3D sampled, simulated, and synthetic datasets. With a booming interest in real-time interactive volume rendering, and with a volume rendering chip being fabricated, volume rendering is a technology whose time has come. In addition to traditional polygonal rendering, commercial graphics systems now offer 3-D volumetric capabilities, and will soon offer volume rendering hardware accelerators.

THE IMPACT OF REAL-TIME VOLUME RENDERING

Visualization of scientific, engineering or biomedical data is a growing field within computer graphics. In many cases the objects or phenomena being studied are volumetric. Volume rendering, which encompasses an array of techniques for displaying images directly from the 3D data, has become a key technology in the visualization of this data.

One of the key advantages of volumetric data is that, unlike surface-based representations, it can embody interior structures and composition. This allows the visualization of internal structure of the 3D data, including amorphous and translucent features. Additionally, the emerging field of volume graphics is using volume data of geometrical models that have been synthesized (voxelized) into 3D datasets, including physics-based modeling of complex objects. Operations such as object cutting, slicing, or tearing, while challenging for surface-based models, can be performed relatively easily with a volumetric representation. More recently, volume rendering has become essential for directly viewing changes of dynamically sampled data, for example, for the visualization of a beating heart under real-time 3D ultrasound.

While advances have been made in the acceptance of volume rendering, two important issues have limited its use:

1) Volume rendering is extremely compute-intensive. Rendering a dataset of 256^3 16-bit voxels at 30 Hz, for example, requires 32 MBytes of storage, a memory transfer rate of 1 GByte per second, and approximately 5 billion instructions per second. This problem is aggravated by the continuing trend towards larger datasets. High-resolution sampling devices, faster supercomputers, and more accurate modeling techniques will make 1024^3 and larger datasets the norm.

2) Volume rendering has complex rendering parameters, particularly transfer functions, which greatly define the appearance of the rendered image. Transfer functions assign values for optical properties, such as color and transparency, to the dataset being visualized. By far the most commonly used method for finding good transfer functions is ``trial and error." Combined with the lack of interactivity, the user often has to find the appropriate parameter settings without immediate visual feedback.

However, lack of real-time frame rates and lack of interactivity are quickly becoming problems of the past. Major advances in desktop computing power, highly optimized volume rendering algorithms, and the emergence of inexpensive hardware for volume rendering will bring volume rendering to a wider audience. Voxel-based games, already sold by several companies, will include new and dazzling volume graphics effects.

This panel discusses the past, present, and future of volume rendering. Presenting a historical perspective and a discussion of their current research, the panelists argue that the main problems of volume rendering are about to be overcome. All panelists represent major corporations, which is an indication that volume rendering has moved from academia to industry. Given the large market potential of volume rendering, the panelists argue that this exciting technology is soon coming to a store next to you.

POSITION STATEMENTS

Arie Kaufman

The high computational requirements of traditional computer graphics led to the development of special-purpose graphics engines, primarily for polygon rendering. Similarly, the special needs of volume rendering, where an image must be computed rapidly and repeatedly from a volume dataset, lends itself to the development of special-purpose volume rendering architectures. A dedicated accelerator, which separates volume rendering from general-purpose computing, seems to be best suited to provide true real-time volume rendering on standard deskside or desktop computers. Volume rendering hardware may also be used to directly view changes of the 3D data over time for 4D (spatial-temporal) visualization, such as in real-time 3D ultrasonography, microtomography, or confocal microscopy. This may lead to the direct integration of volume visualization hardware with real-time acquisition devices, in much the same way as fast signal processing hardware became part of today's scanning devices.

Cube-4, developed at SUNY Stony Brook, is a scalable architecture for true real-time ray-casting of large volumetric datasets. The unique features of Cube-4 are a high bandwidth skewed memory organization, localized and near-neighbor datapaths, and multiple, parallel rendering pipelines with simple processing units. System performance scales linearly with the number of rendering pipelines, limited only by memory access speed. Cube-4 performs arbitrary parallel and perspective projections of high-resolution datasets at true real-time frame rates. The performance is data and classification independent and can be achieved at a fraction of the cost of a multiprocessor computer. Cube-4 uses accurate 3D interpolation and high-quality surface normal estimation without any pre-computation or data duplication. Consequently, Cube-4 is also appropriate for 4D visualization as an embedded volume visualization hardware system in emerging real-time acquisition devices. Possible hardware implementations of Cube-4 for 30 frames per second rates range from an inexpensive PCI board accelerator for

256^3 datasets, to a workstation accelerator board for 512^3 datasets, to a visualization server for 1024^3 or higher resolutions. The cost-performance of Cube-4 is several orders of magnitude better than existing solutions. Both Mitsubishi Electric and Japan radio Co. (JRC) have been fabricating volume rendering chips based on the Cube-4 architecture.

The choice of whether one adopts a general-purpose or a special-purpose solution to volume rendering depends upon the circumstances. If maximum flexibility is required, general-purpose appears to be the best way to proceed. However, an important feature of graphics accelerators is that they are integrated into a much larger environment where software can shape the form of input and output data, thereby providing the additional flexibility that is needed. A good example is the relationship between the needs of conventional computer graphics and special-purpose graphics hardware. Nobody would dispute the necessity for polygon graphics acceleration despite its obvious limitations. We are making the exact same argument for volume rendering architectures.

Marty Brady

The most important single issue in bringing volume rendering into general use is that it is finally available on a general, low-cost computing platform. This has come from a combination of algorithmic advances and steady improvement in computer performance. Volume rendering has also enjoyed somewhat of a free ride on the wave of 3D graphics and multimedia enhancements (e.g., Intel's MMX(tm) technology), which can further accelerate its performance. This rapid pace of advancement will continue. The more difficult next question will be: now that we can do volume rendering, what will we do with it? That is, will new volume rendering applications arise that bring it from a "niche" area into mainstream usage?

Fred Kitson

The transition to volumetric rendering could be characterized as an incremental one over the last decade but that trend seems to be going super-linear as new algorithms and architectures emerge. Applications such as medical visualization have long depended on 3D data collection, display and manipulation, but now more mainstream graphics can benefit from voxel representations. Even game technology for the animation of 3D bodies has been implemented with fast display algorithms that generate compelling imagery. From an API perspective, HP has promoted its "Voxelator" offering as a way to help standardize the application support by OpenGL extensions. From a conceptual level, the approach was to give voxels the same level of support as triangles or pixels. With respect to architecture, this means that one has a volume rendering pipeline as well as a geometry and an image pipeline. This allows integration at the compositing level and engenders support for applications such as image based rendering, polygon modeling of synthesized surface data and volume data. There seems to be an opportunity for volumetric acceleration that exists between the high-end server level and the inadequate instruction level support. Finally, from an applications viewpoint, we see that there is a requirement for mixed rendering such as one would have in the analysis and non-destructive testing of a mechanical part, fluid dynamics, automotive design or surgical simulations. In summary, the time is now for volumetric rendering to escalate in status as part of the mainstream graphics workload and acceleration support.

Bill Lorensen

The landscape of the medical imaging market has changed dramatically over the last few years. At GE we have employed surface extraction and surface based rendering with great success for many applications and many products. Based on our own experience and customer feedback, we concluded that the current market needs are shifting into more direct volume rendering. We thus see 3D volumetric rendering as becoming very important in the near future. We have developed the Visualization Toolkit as a software environment that supports both surface and volume rendering. Our software volume rendering implementations have helped us design abstractions and a visualization pipeline that can accommodate custom hardware-based volume rendering solutions or off-the-shelf workstation and PC-based volume rendering accelerators. We also look forward to hybrid systems that will support combined surface and volume rendering.

Hanspeter Pfister

I will discuss three projects at MERL that address the emergence of the volume graphics market: 1) VolumePro, a PCI card for real-time rendering of 256^3 datasets under WindowsNT, 2) physical-based volume modeling, and 3) Design Galleries for automatic, user guided computation of transfer functions. Mitsubishi Electric will release the VolumePro rendering accelerator in the immediate future. I will also give an outlook on what I believe to be an important research direction, namely the convergence of volume rendering, image-based rendering, and texture mapping towards universal 3D discrete graphics.

BIOGRAPHIES

Arie Kaufman is a Leading Professor of Computer Science and Radiology at the State University of New York at Stony Brook, He is the director of the Center for Visual Computing (CVC) supported by the NSF, DoE, Hughes Aircraft, ONR, NRL, NASA, MERL, HP, SGI, JRC, Howard Hughes Medical Institute, Biotechnology Center, State of New York and many others. His research interests include volume visualization, graphics architectures, algorithms, and languages, user interfaces and VR, and multimedia. He has lectured widely and published numerous technical papers in these areas, including the IEEE tutorial book on Volume Visualization. He has given several courses/papers/panels on Volume Visualization and related areas for SIGGRAPH, Visualization, VBC, CGI, Eurographics and GI. He has been the Papers or Program co-Chair for Visualization '90-'94, co-Chair for several Eurographics/SIGGRAPH Graphics Hardware Workshops, and the chairman of the IEEE CS Technical Committee on Computer Graphics. Kaufman is currently the Editor-in-Chief of the IEEE Transaction on Visualization and Computer Graphics (TVCG). He received a BS in Math and Physics from the Hebrew University of Jerusalem (1969), an MS in Computer Science from the Weizmann Institute of Science, Rehovot (1973), and a PhD in Computer Science from Ben-Gurion University, Israel (1977). He is a Fellow of IEEE, and the recipient of a 1995 IEEE Outstanding Contribution Award and the 1996 IEEE Computer Society's Golden Core Member.

Martin L. Brady received BS, MS, and PhD degrees in Electrical Engineering from the University of Illinois at Urbana-Champaign in 1982, 1984, and 1987, respectively. He has worked as a research scientist at Lockheed's Palo Alto Research Lab, and as an Assistant Professor of Computer Science and Engineering at the Pennsylvania State University. Since 1996 he has been a research scientist in the Microcomputer Research Labs at Intel Corporation in Santa Clara, California. His research interests include volume visualization, computer graphics, parallel algorithms, and image processing.

Bill Lorensen is a Graphics Engineer in the Electronic Systems Laboratory at GE's Corporate Research and Development Center in Schenectady, NY. He has over 25 years of experience in computer graphics and software engineering. Bill is currently working on algorithms for 3D medical graphics and scientific visualization. He is a co-developer of the marching cubes and dividing cubes surface extraction algorithms, two popular isosurface extraction algorithms. Bill is one of the chief architects of LYMB, an object-oriented software development environment written in C. His other interests include computer animation, color graphics systems for data presentation, and object-oriented software tools. Bill is the author or co-author of over 60 technical articles on topics ranging from finite element pre/post processing, 3D medical imaging, computer animation and object-oriented design. He is a co-author of "Object-Oriented Modeling and Design" published by Prentice Hall, 1991. He is also co-author with Will Schroeder and Ken Martin of the book "The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics" published by Prentice Hall in November 1997. He gives frequent tutorials at the annual SIGGRAPH and IEEE Visualization conferences. Bill holds twenty three US Patents on medical and visualization algorithms. In 1991, he was named a Coolidge Fellow, the highest scientific honor at GE's Corporate R&D. Prior to joining GE in 1978, he was a Mathematician at the US Army Benet Weapons Laboratory where he worked on computer graphics software for structural analysis. He has a BS in Mathematics and an MS in Computer Science from Rensselaer Polytechnic Institute.

Frederick Lee Kitson manages the Visual Computing Department at Hewlett Packard's Corporate Research Laboratories in Palo Alto, California. His research efforts include graphics algorithms and architectures, video/audio compression, image processing, and visualization. He is also an adjunct faculty member at the Georgia Institute of Technology. Kitson received a BS degree, with honors, in electrical engineering from the University of Delaware in 1974; an EE degree from the Georgia Institute of Technology in 1975, and a Ph.D. degree in digital signal processing from the University of Colorado in 1981. He was the 1991/1992 Clyde Chair visiting associate professor for the Computer Science Department at the University of Utah.

Hanspeter Pfister is a Research Scientist at MERL— A Mitsubishi Electric Research Laboratory in Cambridge, MA. His research interests include computer graphics, scientific visualization, computer architecture, and VLSI design. He is currently one of the lead architects of VolumePro, Mitsubishi Electric's real-time volume rendering system for PC-class computers. Hanspeter Pfister received his PhD in Computer Science in 1996 from the State University of New York at Stony Brook. In his doctoral research he developed Cube-4, a scalable architecture for real-time volume rendering. He received his Dipl.-Ing. degree in electrical engineering from the Department of Electrical Engineering at the Swiss Federal Institute of Technology (ETH) Zurich in January 1991. He is a member of the ACM, the IEEE Computer Society, and the Eurographics Association.