



The Transfer Function Bake-Off

Organizer:

Hanspeter Pfister, Mitsubishi Electric Research Laboratory (MERL)

Referee:

Bill Lorensen, GE Research

Panelists:

Will Schroeder, Kitware Inc.

Chandrajit Bajaj, University of Texas at Austin

Gordon Kindlmann, University of Utah

Hanspeter Pfister, Mitsubishi Electric Research Laboratory (MERL)

Introduction

As anybody even vaguely familiar with volume rendering knows, finding good transfer functions is very difficult. Pat Hanrahan called it one of the top 10 problems in volume visualization in his inspiring keynote address at the Symposium on Volume Visualization '92. And it seems that today, almost a decade later, there are still no good solutions at hand. Or are there?

In this panel, we pitch four of the currently most promising approaches to transfer function design against each other. The four approaches are:

- The traditional GUI, exemplified by vtk, defended by Will Schroeder, Kitware.
- The Contour Spectrum, proposed by Chandrajit Bajaj, University of Texas at Austin.
- Semi-Automatic Generation of Transfer Functions, invented by Gordon Kindlmann, University of Utah.
- Design Galleries for Volume Graphics, presented by Hanspeter Pfister, MERL.

Each of the four panelists received a couple of volume data sets of unknown origin. Their task was to create meaningful volume visualizations using their respective approaches to transfer function design. In this panel they will present their findings under the scrutiny of Bill Lorensen, chief volume rendering critic and outspoken referee for the panel. Who will be the best? Who will win the grand prize? Or is everybody far off the right track?

Motivation and Panel Format

Direct volume rendering is a key technology for the visualization of large 3D datasets from scientific or medical applications. Of particular importance to the quality of direct volume-rendered images are transfer functions. A transfer function assigns optical properties, such as color and opacity, to original values of the dataset being visualized. Unfortunately, exploring different transfer functions is a tedious task, often done by trial and error. Managing and organizing the exploration of transfer-function space is usually left to the user; the computer is used as a passive instrument. More recent approaches divide the parameter-setting task more evenly between user and computer.

In this panel we will compare and contrast some of the dominant approaches to transfer function design:

- Trial and error, with minimum computer aid (W. Schroeder)

- Data-centric, with no underlying assumed model (C. Bajaj)
- Data-centric, using an underlying data model (G. Kindlmann)
- Image-centric, using organized sampling (H. Pfister)

In order to compare these different approaches we chose the following format:

- Ahead of time, the panelist are given a set of volume datasets of unknown origin or content. They visualize the volume data using volume rendering and their proposed method for transfer function design. The data is provided by Bill Lorensen who is well aware of its content.
- During the panel, each panelist will in turn present a summary of the method in 15 minutes or less, and present the results of the visualization, including:
 - visual results (images and/or animations)
 - performance (timings for each image and memory usage - not including volume rendering)
 - observations (how easy/hard was it, what were the findings, etc.)
- At the end of the panel, Bill Lorensen will discuss the actual content of the volume data, what an experienced visualization practitioner would have hoped to find, and how well the panelists methods achieved this goal. Bill will also announce a winner, if there should be one.
- Bill Lorensen will make sure to keep all presentations short and to the point. This panel and these panelists will not have any trouble generating interesting discussion if there is sufficient time.

This will be a very unique event: how often do alternative approaches to a pressing research problem get to go head-to-head, in person, on real data? This is a genuine scientific experiment, a rarity when evaluating visualization techniques: we have a control (Will Schroeder), we have multiple real-world datasets, and we have an objective quality metric (Bill Lorensen) which we all trust. Also, before visualization we try to get some experience using each other's tools, and thus will have some perspective on their strengths and weakness. We hope that all of this will be done in an atmosphere of lighthearted fun, but with a serious goal, namely to emphasize the importance of further research in transfer function design.

Position Statements

Bill Lorensen: The Illusive Transfer Function

Since its introduction in the late 1980's, Volume Visualization has had limited success in cost effective applications. Advances in image quality and feature sets continue to outpace the technology's acceptance in commercial products.

Several factors contribute to the slow adoption of Volume Visualization:

1. A lack of proven application areas. Routine use of 3D in medicine is still, for the most part, limited to research and teaching hospitals.
2. There is not agreement on software API's for Volume Visualization. This limitation translates to risk for commercial products that adopt one vendor's API over another's.
3. Volume Visualization is slow, requiring expensive workstations with large amounts of memory and special graphics hardware extensions.
4. The Volume techniques are difficult to use by all but an experienced engineer or scientist.

These limitations are being addressed by various universities and companies:

1. New scanners are presenting much more information than a radiologist can possibly review on a slice by slice basis. 3D visualization could be the key to increasing productivity.
2. API's are emerging that fit within current graphics and visualization systems.
3. Low cost, special purpose hardware is now available for personal computers. And, the general purpose processor speeds continue to improve. Texture mapping hardware is available on cheap graphics cards. And, 3D texture mapping could also get cheap if the gamers find a use for it.

However, ease of use is still an issue. Volume visualization has the potential to significantly reduce the amount of time to segment medical data. Fast, robust techniques to create color and opacity transfer functions are needed before volume visualization can move from the lab to the hospital.

Doesn't matter whether the techniques are automatic, semi-automatic or manual. They just need to be fast and simple enough to use.

Will Schroeder

The widespread use of volume rendering has been hampered by the difficulty of creating effective transfer functions. Transfer functions are complex mappings of data (and derived data) values to color and opacity. The complexity of the mapping is further exacerbated by the blending effects along the depth direction. As a result, recent research has focused on automatic and semi-automatic techniques for creating transfer functions.

Such methods are potentially dangerous because the techniques remove the human from the visualization process. Visualization is not just about generating pretty pictures; visualization is also a vehicle of exploration by which the observer comes to understand their data. One can easily imagine semi-automatic and automatic techniques that generate images that fulfill the expectations of the observer, but are not necessarily true to the nature of the data. Thus, it is our contention that transfer functions creation is a necessary

part of the visualization (i.e., data exploration) process. Methods that assist the user in creating transfer functions - and thus improve the efficiency of data exploration - are beneficial. Methods that eliminate the human from the exploration process are dangerous and should be avoided.

Chandrajit Bajaj

In addition to computational and space complexity issues, user interfaces have a tremendous impact on the interactivity of a visualization environment.

A contour spectrum consists of computed metrics over a scalar field. On the basis of such metrics one can define a set of functions which provide a useful tool to enhance the interactive query of the dataset. One primary advantage of the the contour spectrum interface is that allows one to display in a 2D image a "global" view of the examined scalar field, independent of its dimension. For example, in the display of a 3D isosurface, one contour component maybe be hidden inside another. If one associates the isocontour display with the contour tree it becomes immediately clear that the current isosurface is composed of two components and hence one might need a clipping plane to look inside the current isosurface. For time-varying data, functional properties can be computed over time, and displayed using a 2D interface, giving the user a global overview of the time-varying function, allowing interaction in both isovalue and timestep.

I shall report on several examples of contour measures of general utility. The flexibility of the interface allows for numerous enhancements for both general attributes and application specific features.

Gordon Kindlmann

Direct volume rendering is a powerful visualization tool because of its simplicity and flexibility. Both of these characteristics are reflections of the transfer function at the core of the volume rendering process. As simple and "direct" as that mapping is, it is also extremely flexible, because of the immense variety of possible transfer functions. The transfer function's unconstrained flexibility is in fact a big headache. In general, the visualization process should be guided by general information about the goal of the visualization, and specific information about the particular dataset in question. Unfortunately, most interfaces for setting the transfer function are not constrained by either kind of information, so finding a transfer function is usually done through a blind trial and error process.

That this task is frustrating is not surprising when one considers that it is tantamount to performing feature detection by trial and error. Some transfer functions are better than others because they manage to emphasize the features of interest while keeping the rest of the dataset invisible. Most of the time we think of feature detection as happening in the spatial dimensions of the given dataset, but for the purposes of setting a transfer function we need to "locate" features within the space of data values comprising the dataset. It is the lack of a spatial component to this process which makes finding transfer functions frustrating.

My technique addresses this problem for the specific class of volume visualization problems where the goal of the visualization is to see the boundaries between homogeneous regions. It starts with a concept of "boundary" derived from computer vision, but then projects out the spatial component of the boundary by creating a three dimensional histogram of the data value and its first and second derivative. Based on analysis of this histogram, we can create a "distance function" which tries to bridge the unintuitive space of data values to a synthetic, but intuitive, spatial domain: a signed distance to the middle of the nearest boundary. The calculation of the distance function is largely automated. With the distance function in place, the space of transfer functions is usefully constrained

to the visualization goal and the underlying data, making it easier to produce an informative visualization of the boundaries.

Hanspeter Pfister

At the coarsest level of abstraction, all computer-graphics processes map input parameters to output values. For example, a volume-rendering process maps visualization parameters to output pixel values. The trick in producing compelling and useful visualizations is to find input parameters that yield desirable output values. Tweaking input parameters to this end is a tedious experience familiar to anyone who has created computer graphics and scientific visualizations. Design Galleries for Volume Graphics (VolDG) presents a viable alternative to facilitate parameter selection. Instead of asking the computer "What's best?" we ask the computer "What's possible?" The computer's task is to pick a set of input-parameter vectors that spans the space of output values as much as possible; the user's task is simply to select from among the presented possibilities.

VolDG interfaces present the user with the broadest selection - automatically generated and organized - of perceptually different images that can be produced by varying transfer functions. Our current system is built on top of the popular Visualization Toolkit (vtk) and Mitsubishi Real Time Visualization's VolumePro board. The real-time volume-rendering speed of the VolumePro board allows large galleries to be generated in minutes. VolDG is freely available at <http://www.merl.com/projects/dg/>.

The principal technical challenges posed by the VolDG approach are dispersion (finding a set of input-parameter vectors that optimally generates very dissimilar output values) and arrangement (arranging the resulting designs for easy browsing by the user). For dispersion we use a form of evolutionary computation. For arrangement we use multidimensional scaling. The dispersion process can require the rendering of hundreds or thousands of candidate images, and therefore benefits greatly from hardware acceleration by VolumePro. In addition, the interrogative process between user and computer is aided enormously by very expedient rendering.

Biographical Sketches of Panelists

Bill Lorensen

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 30 years of experience in computer graphics and software engineering. Bill is currently working on algorithms for 3D medical graphics and scientific visualization. He co-developed (with Harvey Cline) the marching cubes and dividing cubes surface extraction algorithms, two popular isosurface extraction algorithms. His other interests include computer animation, information display, and object-oriented software tools. Bill is the author or co-author of over 70 technical articles on topics ranging from finite element pre/postprocessing, 3D medical imaging, computer animation and object-oriented design. He is a co-author of "Object-Oriented Modeling and Design" published by Prentice Hall, 1991. He is also co-author with Will Schroeder and Ken Martin of the book "The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics" published by Prentice Hall in November 1997. The text describes an open source visualization C++ class library. He gives frequent tutorials at the annual SIGGRAPH and IEEE Visualization conferences.

Bill holds twenty six 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.

Will Schroeder

Dr. William J. Schroeder is president and co-founder of Kitware, Inc., a small startup company providing commercial support for the open source Visualization Toolkit software. Along with Ken Martin and Bill Lorensen, VTK was initially developed as part of a text on visualization, and is now used around the world in commercial, research, and teaching applications. In his current role at Kitware, Will continues to develop the graphics, visualization, and imaging capabilities of VTK. Kitware is also building proprietary applications on top of VTK, such as the VolView volume rendering system.

Dr. Schroeder received his B.S. in Mechanical Engineering from the University of Maryland, College Park in 1980, and obtained his Ph.D. (while working full-time at the GE Corporate R&D Center) in applied mathematics from Rensselaer Polytechnic Institute in 1991. Will's current research interests include visualization, computational geometry, graphics, and numerical analysis. In 1992 Dr. Schroeder published one of the first papers on polygon decimation. Other areas of interest include swept surface generation, stream polygons for vector field visualization, and object-oriented graphics and visualization architectures. Will is a research faculty member at RPI and continues to explore the use of graphics in the computational sciences.

Chandrajit Bajaj

Chandrajit Bajaj is CAM chair of visualization professor of computer sciences and also current director of the Center for Computational Visualization at the University of Texas at Austin. His research is in the areas of computer graphics, geometric modeling and data visualization. Current approaches include the design and analysis of data structures and compression algorithms that support a blend of image and geometry based multi-resolution approximations. He is also active in developing a new integrated parallel framework for high resolution imaging, finite element simulations, and interrogative visualization.

Gordon Kindlmann

Since 1997, Gordon Kindlmann has been a doctoral student in the Computer Science department at the University of Utah. In 1995 he received a BA in mathematics from Cornell University, and in 1999 he finished his MS in computer graphics under Donald Greenberg in the Program of Computer Graphics at Cornell University. His masters thesis on "Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering" is available at <http://www.cs.utah.edu/gk/MS/>. His current research continues to focus on volume rendering while extending into the areas of medical imaging, segmentation, and color science.

Hanspeter Pfister

Hanspeter Pfister is a Research Scientist at MERL - A Mitsubishi Electric Research Laboratory in Cambridge, MA. He is the chief architect of VolumePro, Mitsubishi Electric's real-time volume rendering system for PC-class computers. His research interests include computer graphics, scientific visualization, computer architecture, and VLSI design. Hanspeter Pfister received his Ph.D. 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 1991. He is a member of the ACM, IEEE, the IEEE Computer Society, and the Eurographics Association.

References

- [1] C. Bajaj, V. Pascucci, and D. Schikore. The Contour Spectrum. In *Proceedings of the 1997 IEEE Visualization Conference*, pages 167–173. Phoenix, AZ, October 1997. Online at <http://www.ticam.utexas.edu/CCV/papers/fhighlights.html>.
- [2] G. Kindlmann and J. Durkin. Semi-Automatic Generation of Transfer Functions for Direct Volume Rendering. In *Proceedings 1998 IEEE Symposium on Volume Visualization*, pages 79–86. 1998. Online at <http://www.cs.utah.edu/gk/papers/vv98/>.
- [3] J. Marks, B. Andalman, P.A. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. Design Galleries: A General Approach to Setting Parameters for Computer Graphics and Animation. In *Proceedings of Siggraph 97*, pages 389–400. ACM Siggraph, Los Angeles, CA, August 1997. Online at <http://www.merl.com/reports/TR97-14/>.