The Future Visualization Platform

Panel Organizer:
Greg Johnson, University of Texas at Austin

Panelists:
David Ebert, Purdue University
Chuck Hansen, University of Utah
David Kirk, NVIDIA Corporation
Bill Mark, University of Texas at Austin
Hanspeter Pfister, Mitsubishi Electric Research Laboratories

INTRODUCTION

Advances in graphics hardware and rendering methods are shaping the future of visualization. For example, programmable graphics processors are redefining the traditional visualization cycle. In some cases it is now possible to run the computational simulation and associated visualization side-by-side on the same chip. Moreover, global illumination and non-photorealistic effects promise to deliver imagery which enables greater insight into high resolution, multivariate, and higher-dimensional data.

The panelists will offer distinct viewpoints on the direction of future graphics hardware and its potential impact on visualization, and on the nature of advanced visualization-related tools and techniques. Presentation of these viewpoints will be followed by audience participation in the form of a question and answer period moderated by the panel organizer.

Keywords: future, visualization, hardware, techniques

POSITION STATEMENTS

David Ebert:
Effective Image Generation Techniques for Visualization

The value of and evaluation criteria for visualization must be effectiveness and usability, not the rendering technique that is used. The increase in performance and programmability of PC graphics hardware has enabled complex volume rendering techniques to perform at interactive rates on moderate sized datasets. Therefore, many techniques can be chosen to create visualizations from datasets. The issue now becomes which techniques should be used to produce an effective and controllable visualization. Do we want to use 15 dimensional transfer functions with 100 parameters? Do we want multiple scattering, wavelength-dependent illumination and par-

ticle scattering of the dataset? Will a simple sketch of the important structures and objects of interest suffice? Or will advanced analysis and a simple text answer be more effective?

We have been developing rendering techniques ranging from physics-based multiple scattering illumination models to nonphotorealistic illustrative visualization of datasets, all with the goal of providing users with the information they seek from their data. We have also been developing techniques to abstract the controls of these systems to a more intuitive, user and domain oriented interface to increase the usability and reliability.

Chuck Hansen:
Global Illumination for Visualization

Interaction with complex, multi-dimensional data is now recognized as a critical analysis component in many areas, including computational fluid dynamics (CFD), combustion modeling and simulation, and medical simulation and imaging. Scientific and engineering applications such as these are increasingly utilizing teraflop-class parallel computers which offer enormous potential for solving very large-scale problems. However, making effective use of this potential relies upon the ability of human experts to interact with their computations and extract useful information from the resulting 3D volumetric data sets.

Direct volume rendering has proven to be an effective and flexible visualization method for interactive exploration and analysis of 3D scalar fields. While widely used, most if not all applications render (semi-transparent) surfaces lit by an approximation to the Phong local surface shading model. This model renders translucent surfaces simplistically and does not provide sufficient lighting information for good spatial acuity. In fact, the constant ambient term leads to misperception of information that limits the effectiveness of visualizations. Furthermore, the Phong shading model was developed for surfaces, not volumes. The model does not work well for volumetric media where sub-surface scattering dominates the visual appearance (e.g. tissue, bone, marble, and

atmospheric phenomena). As a result, it is easy to miss interesting phenomena during data exploration and analysis. Worse, these types of materials occur often in modeling and simulation of the physical world.

Physically correct lighting has been studied in the context of computer graphics where it has been shown that the transport of light is computationally expensive for even simple scenes. Yet, for visualization, interactivity is necessary for effective understanding of the underlying data.

We seek increased insight into volumetric data through the use of more faithful rendering methods that take into consideration the interaction of light with the volume itself. The future visualization platform will include new robust interactive volume shading methods that incorporate global illumination effects.

David Kirk:
The Evolution of Graphics Processors

Graphics processor architectures have evolved very rapidly in the past few years. After some years of simply implementing and accelerating the OpenGL and DirectX graphics pipelines, the innovation in GPU development has moved on to a different approach. While GPUs are still providing a capability to accelerate the graphics APIs, the formerly hardwired and hardcoded functionality of GPUs is being abstracted into more general purpose programmability. GPUs are becoming highly parallel programmable streaming floating point engines, and for a given amount of silicon area, GPUs are far more powerful than other more general computing devices such as CPUs.

The reason for this difference lies in the way that GPU architectures have evolved, and the fact that this evolution has taken a very different path than CPUs. As GPUs become more flexible, powerful, and programmable, their architecture is more well-suited to embrace the parallelism that is inherent in graphics, shading, and other hard computational problems. In order to harness this compute power for non-traditional graphics problems, the programming and dataflow must be adapted to the GPU architecture. Large computational problems which can be organized as data streaming through the GPU can be accelerated dramatically. This can have profound effects on data visualization.

Bill Mark:
The Return to Generalized Hardware –
A Visualization Renaissance

The addition of user-programmable functionality to graphics hardware has already had a significant impact on the visualization community, but so far we've only seen a small piece of the potential impact of this transformation. Within the next 4-8 years, we will cease to talk about "graphics hardware", and will instead talk about "parallel hardware" that happens to be good at executing parallelized real-time rendering software. This "parallel hardware" will be sufficiently

flexible that it will also perform well for a variety of other parallelizable tasks such as simulation computations.

The visualization community will benefit even more from this transformation than the mainstream entertainment industry will. Once most rendering computations are expressed as high-performance parallel software rather than hard-wired algorithms, it will be easy to deploy rendering algorithms specialized for niche markets such as visualization, even if those algorithms are somewhat different from those used by the entertainment industry. For example, it will be possible to use global illumination algorithms rather than Z-buffer algorithms, or to efficiently visualize custom-compressed volume data sets rather than raw volume data. Even more importantly, it will be possible to tightly integrate simulation with visualization on a single parallel computation engine.

Hanspeter Pfister:
Advanced Tools and Techniques – A Shift in Focus

The future of visualization does not depend on what particular graphics hardware we will be using. Rendering is – for all practical purposes – a solved issue. The future of visualization is building visual tools to aid in the human-guided analysis of large, possibly time-varying data sets. To be successful we need to draw from experiences in human-computer interaction, vision, perception, machine learning, and visualization. A future visualization tool needs to present information in a concise and useful format to enable the user to interpret the data; provide a rich and fluid interaction with the data to facilitate exploration and discovery; and must contain powerful learning algorithms to expose patterns that would be inaccessible to the human given raw data alone. I will briefly present a few MERL projects that touch on these core requirements.

- **DiamondTouch** is a simultaneous, multi-user, touch sensitive input device developed at MERL. Not only can it detect multiple, simultaneous touch events, but it can also identify which user is touching where.

- **Multi-Parametric Visualization** is a set of methods and tools for visualization and querying of multidimensional information. It is designed to be easy to use and quick to yield insights into the data.

- **Incremental Imputative Singular Value Decomposition (IISVD)** allows an SVD to be computed from streaming data. The technology is distinguished both by its speed – it is the first linear-time single-pass algorithm – and by its ability to handle tables with many missing elements – a common problem in data mining.

- **Computer-Human Observation** is a large internal project, currently in its third year, with the goal of defining how computer vision and learning technology will shape the future of visual surveillance. One particularly relevant aspect of this umbrella project is the focus on

employing learning and image analysis technology to help analysts make sense of enormous visual surveillance databases.

BIOGRAPHICAL SKETCHES

David Ebert

David Ebert is an Associate Professor in the School of ECE at Purdue University and received his Ph.D. from the Computer and Information Science Department at The Ohio State University in 1991. His research interests are scientific, medical, and information visualization, computer graphics, animation, and procedural techniques. Dr. Ebert performs research in volume rendering, illustrative visualization, minimally immersive visualization, realistic rendering, procedural texturing, modeling, and animation, and modeling natural phenomena. Ebert has been very active in the graphics community, teaching courses, presenting papers, chairing the ACM SIGGRAPH 97 Sketches program, co-chairing the IEEE Visualization '98 and '99 Papers program, serving on the ACM SIGGRAPH Executive Committee and serving as Editor-in-Chief for IEEE Transactions on Visualization and Computer Graphics. Ebert is also editor and co-author of the seminal text on procedural techniques in computer graphics, Texturing and Modeling: A Procedural Approach, whose third edition was published in December 2003.

Chuck Hansen

Charles Hansen is an Associate Professor of Computer Science in the School of Computing, and the Associate Director of the Scientific Computing and Imaging Institute at the University of Utah. His research interests include large-scale scientific visualization and computer graphics. He has been an active contributor to, and organizer of, the IEEE Visualization Conference. He is currently Associate Editor in Chief of the IEEE Transactions on Visualization and Computer Graphics. He received a B.S. in computer science from Memphis State University in 1981 and a Ph.D. in computer science from the University of Utah in 1987. From 1997 to 1998, he was a Research Associate Professor in CS at Utah. From 1989 to 1997, he was a Technical Staff Member in the Advanced Computing Laboratory (ACL) located at Los Alamos National Laboratory, where he formed and directed the visualization efforts in the ACL. He was a Bourse de Chateaubriand PostDoc Fellow at INRIA, Rocquencourt France, in 1987 and 1988.

David Kirk

David Kirk is Chief Scientist and Vice President of Architecture at NVIDIA. He was previously Chief Scientist and head of technology for Crystal Dynamics, and prior to that worked on developing graphics hardware for engineering workstations at Apollo/Hewlett-Packard. David holds B.S. and M.S. degrees from MIT and M.S. and Ph.D. degrees from the California Institute of Technology, and is the author/inventor of over 100 technical publications and patents in the area of computer graphics and hardware. At SIGGRAPH 2002 in San Antonio, TX, David was the recipient of the ACM SIGGRAPH Computer Graphics Achievement Award, honoring him for his contributions to the field.

Bill Mark

William R. Mark has worked on programmable graphics systems in both academia and in industry. From 2001 to 2002, he worked at NVIDIA as the lead designer of the Cg language, a programming language for graphics hardware. This project was an industrial follow-on to his earlier postdoctoral research on similar systems at Stanford University with Kekoa Proudfoot, Pat Hanrahan and others. Bill has now returned to academia as an Assistant Professor at the University of Texas at Austin, where he and his research group are investigating algorithms and architectures for future real-time graphics/parallel systems. Bill received his B.S. degree in physics from Rice University in 1992 and his Ph.D. degree in computer science from the University of North Carolina, Chapel Hill in 1999. Last year he served as the papers co-chair for the SIGGRAPH/Eurographics 2003 Conference on Graphics Hardware.

Hanspeter Pfister

Hanspeter Pfister is Associate Director and Senior Research Scientist at MERL - Mitsubishi Electric Research Laboratories - in Cambridge, MA. He is the chief architect of VolumePro, Mitsubishi Electric's real-time volume rendering hardware for PCs. His research interests include computer graphics, computer vision, scientific visualization, and graphics architectures. His work spans a range of topics, including point-based graphics, 3D photography, 3D television, face modeling, face recognition, and volume graphics. Hanspeter Pfister received his Ph.D. in Computer Science in 1996 from the State University of New York at Stony Brook. He received his M.S. in Electrical Engineering from the Swiss Federal Institute of Technology (ETH) Zurich, Switzerland, in 1991. Dr. Pfister has taught courses at major graphics conferences including SIGGRAPH, IEEE Visualization, and Eurographics. He has been teaching introductory and advanced graphics courses at the Harvard Extension School since 1999. He is Associate Editor of the IEEE Transactions on Visualization and Computer Graphics (TVCG), chair of the IEEE Technical Committee on Visualization and Computer Graphics (TCVG), and has served as a member of international program committees of major graphics conferences. Dr. Pfister was the general chair of the IEEE Visualization 2002 conference. He is senior member of the IEEE, and member of ACM, ACM SIGGRAPH, the IEEE Computer Society, and the Eurographics Association.