Report on the
DOE Workshop on Tera-Scale Visualization

Lawrence Livermore National Laboratory

Livermore, California

September 9-12, 1997

Edited by Roger Crawfis

Executive Summary

This report outlines the discussions and results of a small workshop held at Lawrence Livermore National Laboratory (LLNL) on the visualization and comprehension of large terabyte data sets. Particularly those data sets generated from computational physics and engineering using modern massively parallel supercomputers. The format of the workshop consisted of one day in which various physicists and engineers discussed the use of the new massively parallel architectures, the data complexities being generated, and their requirements for managing and visualizing this data. The remaining three days were spent discussing various technologies that could be applied to the problems, discussing their relative strengths and weaknesses, and assessing their potential impact. The technologies were roughly grouped into three overlapping areas:

- Data Reduction Algorithms
- Data Comprehension techniques
- Interactive Technologies
- Architectural and System Design

Nearly fifty different technologies were identified as potentially being applied to the problem. Due to time constraints, a small sampling of these technologies were further discussed, and a subjective assessment of their potential impact, the uncertainty (or risk) of that impact and an estimated required effort for these technologies was determined. Overall, those techniques that bring us directly closer to the end-users inquiry were determined to have the most potential impact. These included automatic feature extraction, data mining, and comparative analysis. It was felt that some progress can be made on feature extraction with a minimal effort, but substantial progress exhibited a higher degree of uncertainty as to whether the expended effort would lead to a very large impact. It was also felt, that all of the technologies discussed would greatly aid in the visualization effort, and to ensure success, parallel efforts should be undertaken in many of these technologies.

Technologies

Most of the workshop focused on determining technologies that might best bear fruit when applied to the problems of tera-scale visualization. This section briefly outlines each of these technologies in that context. Several of these technologies were further examined for their potential impact versus risk and investment. This impact will be elaborated on in the next section. Due to the limited time frame of the workshop, only a small fraction of technologies was discussed in this vein. The grouping of the technologies placed several technologies in more than one category. These are indicated in Table 1.

Data Reduction	Data Comprehension	Interactive Techniques	Architectural / System
Hierarchical and Multi-resolution	Data Analysis	View-dependent Acceleration	Software Engineering / Architectures

Techniques		Techniques	
Paging / Caching	Feature Extraction / Tracking	Distributed / Remote Visualization	Paging / Caching
Pre-processing for Interaction	Data Mining	Virtual Reality	Parallel Algorithms
Polygon Reduction Algorithms	Comparative Techniques	Image-based Rendering	Data Models
Higher-order Basis Functions	Uncertainty and Error Visualization	New Hardware Architectures	Scientific Data Management
Feature Extraction / Tracking	Visual Metaphors	Pre-processing for Interaction	Program Visualization
Data Server	Derived Variables		Distributed Visualization
View-dependent Algorithms	Experimental Visualization / Acquisition		Pre-planning Tools
Data Compression	Time-varying Analysis		Remote Visualization
	Sensitivity Visualization		Error Analysis
	Simulated Experimental Cameras / Devices		End-to-end System Architectures

(Technologies in alphabetical order)

Architectures / Remote Visualization

The display of the visualization results is rarely directly on the parallel machine (a HIPPI frame-buffer may be the rare exception). The results must be displayed down to the users desktop, or less preferred, to a near-by office dedicated for this purpose. Given a specific architecture, what is the appropriate partitioning of the workload for the data management, transformations, mappings to visual entities and the display of the resulting objects across the architecture? If networks are a bottleneck, can we compress the data or only transmit differences from the previous view or contour? Can we leverage any graphics hardware acceleration or data servers? Is quality of the image more important than interactivity for the end-user? An efficient and effective system will need to examine all of these issues to select an appropriate architecture. Even after the architecture is selected though, much development and research is needed to maximize the throughput, and ameliorate the bottlenecks.

Impact

Unknown.

Previous Work

Unknown.

See Also: Data Server, End-to-end System Architectures

Comparative techniques

A single simulation is usually part of a larger study, composed of many computer simulations and experiments. The final task is usually not to understand a single simulation run, but to understand that simulation and how it relates to previous simulations or actual experiments. The end-users thus have a need to:

- Compare simulations from the same code system using different parameters or initial conditions.
- Compare simulations from one code system to those of another system.
- Compare computer simulations to results obtained from experimental simulations.
- Compare different experimental simulations.
- Compare the state of a simulation from one time to another time.

Code developers must likewise compare many different computer simulations to each other to verify modifications to the code are correct. This merging of many disparate data sources is a fundamental problem driving ASCI and many other computational science efforts. Can we develop and verify computer codes which can accurately represent experimental test data, and then apply these codes to new virtual experiments with a high degree of confidence in their results.

Very few visualization tools handle multiple data sets or even multiple time steps very effectively. Research into better techniques for managing and representing multiple data sets are needed. These techniques need to represent the data both qualitatively and quantitatively, handle different grid types, ensuring proper resampling, interpolation and registration, and present the resulting multi-variate information comprehensibly.

Impact

As mentioned, this is the driving force of DOE's ASCI program. The impact is thus extremely high, and should be a target area for the ASCI Visualization program.

Previous Work

While some visualization systems allow merging of individual images and some basic two-dimensional comparison, very little research and development has gone into addressing this problem directly and completely.

See Also: Experimental Visualization / Acquisition, Visual Metaphors

Data analysis

This category, also termed Stats+, includes Data Mining. It differs from Feature Tracking in that there is not a "thing" that the algorithms are directly searching for. The goal is to generate various meta-data, which requires much less memory and storage. Examination of this meta-data would then hopefully lead to insight or aid in setting the various visualization parameters (e.g., determining a time-step to examine, etc.).

Possible activities in this area include:

- Determining outliers in the data
- Generating histograms in many new ways for different properties and data slices
- Determining and displaying various statistical properties over a selected region (ala. a magic mirror)
- Providing hierarchical representations of the properties, using various "averaging" techniques.
- Applying various correlation techniques
- Generating many different data mappings (or transforms)
- Automatically determining useful transforms.
- Providing information for efficient caching and data paging.

Impact

The impact for this area is unknown, but potentially high.

Previous Work

Histogram techniques have been around for years and have proven useful for image processing. [Bajaj97], has recently shown the use of histograms of precomputed properties to aid in iso-contour selection.

See Also: Pre-processing for Interaction

Data Compression

The shear volume of data is difficult or impossible to transport and manage. Lossy and lossless compression techniques of the raw data need to be investigated. This would aid not only the disk storage of the data, but also the retrieval of the data, the resulting computation of derived variables (assuming an equivalent accuracy can be guaranteed), and with progressive compression techniques, the transmittal and possibly selection of the data.

Impact
Unknown.
Previous Work
Previous work in this area has focused mainly on regular or curvilinear data where coherency between neighbors can easily be utilized. Octrees, BSP-trees and k-d trees have been studied for many years, while wavelet decompositions have shown recent interest. Internal experiments within LLNL have shown compression ratios of 25:1 for octrees and close to 100:1 for run-length encoding with a ten- percent loss in accuracy. Difference encodings, both spatially and over time, may yield better compression ratios. These are typically employed in wavelets and with MPEG video compression. LLNL is currently looking into using MPEG compression for two-dimensional slices of a 3D time-varying data set. Most of these techniques require a uniformly discritized data range that is mapped to a 8 to 32-bit integer. This can be problematic for data fields with a wide-varying range. Previous research has examined a separate lossless compression of the exponents, and a lossy compression of the mantissa.

See Also: Higher-order Basis or Interpolation Functions, Data Server

Data Mining
(see Data Analysis, Feature Extraction / Tracking)

Data models/scientific data management
Different simulation packages utilize different data or mesh structures. These different samplings require substantially different visualization algorithms for the same basic technique. Providing a full set of visualization tools for every possible data topology and every possible data structure encoding that data topology requires a huge investment in development. Abstracting these data topologies to several key data structures and basic operations is needed to mitigate this workload explosion. Alternatively, mechanisms for promoting or demoting different data types to those needed for the visualization algorithms may be developed. This may lead to inefficiencies in the algorithms or data storage (for the case of promotion), or concerns over accuracy (for the case of demotion).
Impact
Unknown.
Previous Work
Unknown.

See Also: Image-based Rendering, Data Server

Data Server: Selection & Reduction
Commercial servers are rapidly being developed optimized for fast I/O access and delivery. These machines work well for selecting and transmitting data. With better user interfaces and algorithms and more development, the user sitting at a desktop can select a smaller region of data with a desired resolution using this new breed of computers. A request is then sent to the Data Server to select and transmit this modest data for visualization on the desktop. More work is need on a richer and fuller set of queries.
Impact
Unknown.
Previous Work
Unknown.

See Also: Remote Visualization

Derived Variables
Many simulations calculate and store as few variables as possible. Variables derived from the stored data may be more meaningful for the user’s analysis. Calculating these derived variables must be dealt with in any visualization system, and can represent a substantial portion of the system in terms of both lines of

programming and computational time. With very large data sets, this can be an ardent undertaking. New research into performing lazy evaluations of the variables, with efficient caching and partial representation is needed in this area.

Impact

This is an area that is done and must continue to be supported. Whether any additional research into this area is required is unknown. Systems such as AVS and Data Explorer easily support extending the system for application specific derived variables.

Previous Work

Unknown. Application specific.

See Also:

Distributed Visualization – heterogeneity

(see Architectures / Remote Visualization)

End-to-end System Architectures

Seven possible high-level architectures that are currently used or being considered were briefly identified. These were as follows:

1) A frame-buffer connected directly to the MPP via a HIPPI bus.

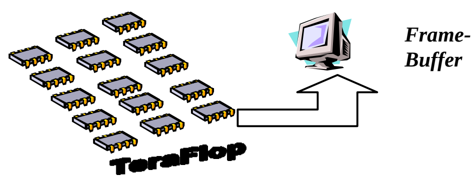

2) A Data Silo is used to store the data from the MPP and sub-sequentially accessed by the visualization system.

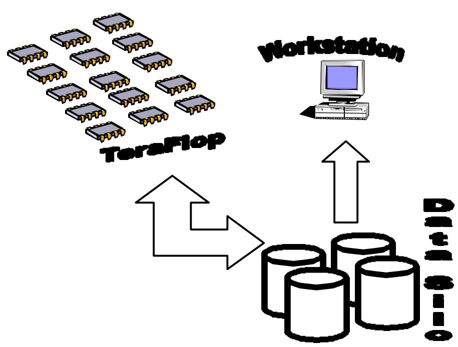

3) A high-end visualization server is attached to the MPP and also directly to a data silo. The imagery is optionally exported to a desktop workstation.

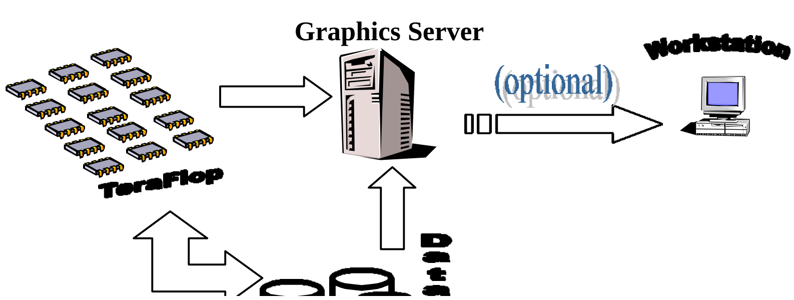

4) The MPP includes specialized high-speed graphics hardware, which is used to generate the imagery for display on a workstation.

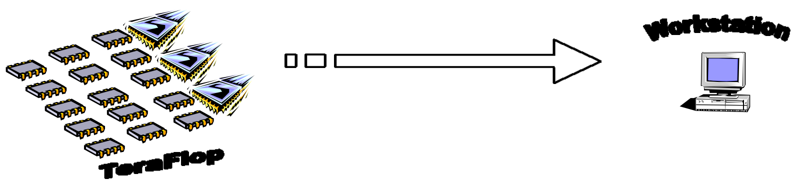

5) Similar to #3, but with the Graphics Server replaced by a Data Server. The workstation is no longer optional.

6) A future configuration might have an integrated MPP and data server(DS), which access a data silo. A Graphics server and individual workstations are tied directly to the data server portion of the MPP/DS.

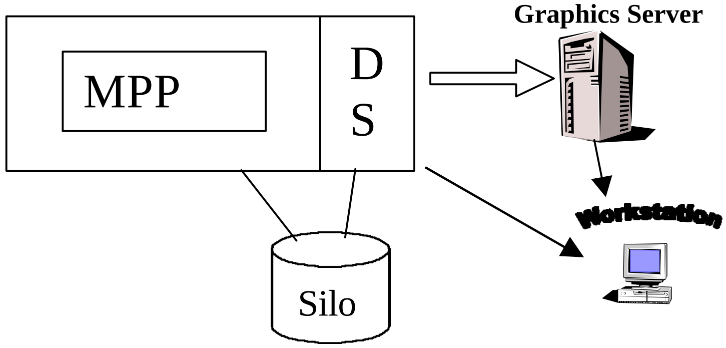

7) Two MPP's are incorporated in the final future configuration. One is oriented towards batch processing, and has a very high-speed connection to a more interactive MPP or tera-flop machine. This interactive machine is connected to the users' desktop.

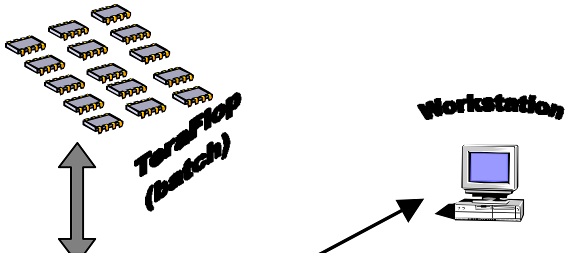

Experimental Visualization / Acquisition

Of primary importance to the end-users, is the ability to understand experimental data. Extracting information from legacy pictures, in which no further constraints can be placed on the environment, is required for ASCI. Some experiments are still conducted to aid in verifying the simulation codes. A tighter integration of visualization or signal reconstruction technologies with the acquisition devices, the registration and the placement of these devices are needed to ensure accurate quantitative information can be extracted from the resulting images. Technologies from pattern recognition, computer vision, three-dimensional image reconstruction and stereo image analysis should be incorporated into the complete visualization processes.

Impact

Unknown.

Previous Work

Whole disciplines have active research into 3D reconstruction from images, image registration, range cameras and image processing. Much of this technology needs to be evaluated and incorporated into this area.

See Also: Comparative Techniques, Visual Metaphors

Feature Extraction / Tracking

Often, the physicist or engineer is interested in finding new features or anomalies within their dataset. This may include vortices or vortex cores in a fluid flow simulation or frontal systems for a weather simulation. There are many different technologies that can come into play on this such as machine vision or expert systems. Once features are detected, we have the added complexity of following their evolution over time. In many cases, features can be easily expressed. The threshold or iso-value of a function may represent a point of failure. Often a visualization will convey the feature visually enabling the user to easily identify the feature on the image.

Impact

Extremely High.

Previous Work

Some work has been done in the area of fluid flow and vector field analysis.

Hierarchical and Multi-resolution Techniques

In an effort to maintain interactive rates for the rendering of the visualization pipeline, many researchers are examining hierarchical representations. These data structure store the geometric information at various levels of detail or LOD's. A specific LOD can then be selected to maintain a constant frame rate during the rendering process. When the user pauses from the navigation, progressive refinement techniques can be used to update the display with a higher level of detail. This maintains interactivity while navigating, while providing accuracy during idle cycles. Research into multi-resolution representations for the raw data is needed to allow for interactive visualization queries, more efficient feature tracking algorithms, and many other visualization algorithms. Furthermore, research into acceptable LOD's and error indicators is needed to ensure the user does not navigate past critical data points. Multi-resolution techniques that can retain significant features, such as discontinuities, are needed to ensure the navigation is meaningful. Currently, most, if not all, of the current research for volumetric data is being conducted for regular or curvilinear grid topologies. Research into structures and techniques that operate and maintain the underlying finite-element grid structures is needed.

Impact

Potentially High.

Previous Work

Quadtrees, Octrees, Wavelets, Progressive Meshes.

See Also: View-Dependent Acceleration Techniques, Polygon Reduction

Higher-order Basis or Interpolation Functions

Most visualization algorithms utilize a tri-linear interpolation scheme within individual mesh elements. More accurate interpolation algorithms would reduce the errors introduced by these methods. More importantly, the use of higher order basis functions can be used to provide a reconstruction function that may be stored more succinctly.

Impact

Unknown.

Previous Work

Unknown.

See Also:

Image-based Rendering

For very large complex models, efficiency gains can be achieved by precomputing rendered images for many different views and then selecting the appropriate image as the user navigates through the model space. This decreases the non-productive time of the user by reducing the latency time of delivered information (images). Image formats lend themselves easily to compression and multi-resolution formats. Furthermore, many companies are working on strategies for streaming high-resolution video across the network to a low-cost set-top device. In addition to visual images, we can store in image format a rasterized version of slices through the data, as well as store the z-buffer and other information for delivery-end video merging, compositing, and manipulation. Research is needed into resampling issues for representing quantitative information, and quantifying any errors introduced in the transformations. Determining the appropriate set of views, data variables and contouring or volume rendering parameters also needs to be explored.

Impact

Unknown.

Previous Work

QuicktimeVR, Lumigraph, Light Fields, Plenoptic Rendering, LDI’s.

See Also: Hierarchical Methods, Visual Metaphors

Paging / Caching

Performance in visualization is dictated by many factors. In addition to the obvious calculational and rendering costs, the access of data can be substantially high. Many visualization techniques only require a small fraction of the total data set. A streamline calculation, iso-contour surface generation or simple slice plane of the data typically require less than 10% of the total data. Requesting a slice or particle path through the time domain can also be problematic to most visualization systems. Intelligent paging and caching of the data can more readily allow flexibility in the access schemes.

Impact

Unknown.

Previous Work

Unknown.

See Also:

Parallel Visualization Algorithms

Since the data is being generated on the parallel supercomputers, some components of the visualization process will certainly require parallel implementations on these systems. Previous work in this area has focused on parallel rendering, and re-distributing the data for optimal rendering performance. More research is needed into algorithms utilizing the domain decomposition of the underlying simulation. Harder research problems that the workshop attendees identified include, polygon reduction techniques, the rendering process, data compression, feature extraction, data mining and view dependent algorithms. The resulting algorithms are greatly tied to and influenced by the underlying visualization architecture.

Impact

High. The techniques will be pervasive, as many of the other technologies will need to perform their work in this domain if the data can not be migrated to a single CPU machine.

Previous Work

The ACM has a Symposium devoted to this subject, and it has been an active area of research for both academia and the national labs.

See Also:

Polygon Reduction Algorithms

Many visualization algorithms convert the raw data to a geometric description for rendering. Reducing the complexity of this geometric representation can greatly enhance the continual rendering of the geometry as the user interacts with the system. This category encompasses techniques that take a triangular (or polygonal) mesh and produce a separate triangular mesh having fewer triangles, but deviating from the original mesh by a small error. Some of these techniques have already been incorporated into existing visualization systems. Parallel implementations exist for some of the techniques, but more research is needed in this area. An iso-contour surface may be color-coded to represent an additional data variable. As such, the polygon reduction must take into account the mesh attributes in addition to the mesh shape and topology. Future research is needed in volume (or polyhedra) reduction techniques and hierarchical representations.

Impact

Unknown.

Previous Work

This is a very active area of current research. The techniques have widespread applicability in GIS systems, terrain modeling, scientific visualization, animation, virtual environments and game development.

See Also: Hierarchical Representations

Pre-planning Tools

For several of the technologies listed -- Pre-processing for Interaction, Data Mining, Image-based Rendering, Data Analysis, Feature Tracking, Data Servers and possibly many others -- the specification of which data to manipulate, rearrange, etc., and how it should be performed can be problematic. Where should histograms be calculated, What iso-contours should be precalculated, and what set of images should be generated and for what views, are just some of the questions that need to be asked as well as specified before the start of the simulation. As more and more of these questions need to be pre-specified, more powerful and flexible tools need to be developed to aid the user in this specification. This may either be in an interactive environment with a previous small-scale simulation of the experiment, or using intelligent algorithms to automatically detect what it should do (based either on the data or on the user's past tendencies).

Impact

Unknown.

Previous Work

Unknown.

See Also: Image-based Rendering, Pre-processing for Interaction, Data Analysis

Pre-processing for Interaction

The visualization pipeline consists of many different steps before a final number or image is produced. Pre-computing several of these steps allows a quicker response back to the user. The area of Image-based rendering is one technology in this category, that includes most (if not all) of the visualization pipeline. An intermediate step, might be to precompute iso-contour surfaces or geometry display lists for later rendering. Polygon reduction techniques may even be applied to these as a pre-process. Other areas for possible inclusion are the Data Analysis techniques, search structures such as interval trees and histograms, and file structure optimizations.

Impact

Unknown.

Previous Work

Unknown.

See Also: Pre-planning tools, Image-based Rendering

Program Visualization

The experiment is only as good as the equipment controlling (or diagnosing) the experiment. The development of robust codes is critical to the ASCI efforts. Analyzing the program's behavior, rather than simply its output, may prove more beneficial than current techniques.

Impact

Unknown.

Previous Work

Unknown.

See Also:

Sensitivity Visualization

Many visualization techniques show a coarse sample point or a thresholding of the data. Moving to a nearby point may produce dramatically different streamlines in an advection-based algorithm. Likewise, changing the iso-value for a contour surface generator may dramatically change the shape and topology of the resulting surface. This sensitivity to the user's parameters needs to be better understood and conveyed to the user to avoid misconceptions about the data.

Impact

Unknown.

Previous Work

Unknown.

See Also: Uncertainty and Error Visualization, Visual Metaphors

Simulated Experimental Cameras/Devices

Experimental scientists are accustomed to looking at the imagery or output of specialized cameras and other devices applied to real experiments. For virtual experiments, visualization that mimic these devices allow a familiar view for the scientists. More importantly, they allow for a better comparison between actual experiments and the virtual experiments. New, application-specific techniques are needed to mimic these devices.

Impact

Unknown.

Previous Work

Unknown.

See Also: Experimental Visualization / Acquisitions, Visual Metaphors

Software Engineering / Architectures

Extensibility, maintainability, and reliability are necessary for any sustained productivity. If various aspects of a visualization environment can be abstracted for reuse, more researchers can examine the problems of ASCI visualization without a large expenditure in the basic infrastructure - organizing, reading, manipulating the data files, dealing with specific grid topologies, accessing derived variables, etc. Frameworks such as IBM's DX and AVS's Express allow for simplified extensibility on the data manipulation and calculations. SGI's Open Inventor allows for easily extension of the rendering techniques.

Common data formats, API's and specified software practices would greatly increase the code reuse.

Impact

Unknown.

Previous Work

Unknown.

See Also:

Time-varying Analysis

In an idealized world, we may only be concerned with where we are going, rather than with how we got there. In order to detect deficiencies in the codes or deviations from legacy data, an examination of the simulation at many time-steps is important. We can quite often learn more about the evolution of the simulation, by including the time domain in the visualization. This may be achieved by changing the visualization over time, or by remapping the visualization tools to treat time as a full-fledge coordinate axis. Most tools are not set-up to handle time as the latter. The data organization and management problems are compounded greatly. Research into new data formats, paging schemes and caching is needed to allow better access to slices of the data along the time axis. For this to be useful, adequate sampling of the data in time is also needed. This is often not the case with most simulation runs, due to limited disk space and poor data management. Finally, new visual representations that illustrate the evolving nature of a simulation clearly are needed. Compression of the data in the time domain may alleviate the data size complexities.

Impact

Unknown.

Previous Work

UFAT by NASA/Ames pre-computes particle traces in unsteady flows and allows for rapid playback of these data sets. LLNL has examined slices in the x/y/time domain that were then JPEG compressed.

See Also: Comparative Analysis, Visual Metaphors

Uncertainty and Error Visualization

With any visualization technique a user uses, there is a certain amount of possible error. The user needs to have a measure for these errors. This may either be visually or by listing various assumptions (such as, bilinear interpolation across zones, etc.), global error properties and other meta-data (the steps and data mappings undertaken to arrive at the visualization). Coupling the error in the physics simulations with the visualization error bounds is needed to show the compound errors. For a visual representation, new visual metaphors are needed to represent the errors in the context of the underlying data field.

Impact

Unknown.

Previous Work

Unknown.

See Also: Sensitivity Visualization

View-dependent Acceleration techniques

Many polygons in large complex scenes or data sets, do not contribute to the final image a user is currently viewing. This may be due to occlusion. There are other polygons closer to the user’s viewpoint that are hiding polygons behind it. It may also be due to the limited field-of-view of the computer monitor. Polygons to the side of the user are not displayed with conventional display devices. Finally, the model may be so complex, that many polygons fall on the same pixel and are contained within that pixel. This may lead to aliasing of even missing small data features contained in these tiny polygons. The section on Hierarchical and Multi-resolution Techniques will cover the research needed for this later category. For the other cases, research is needed into efficiently locating and managing occluders and view frustum clipping. Several research papers have appeared on these topics over the past few years. Products, such as SGI’s new OpenGL Optimizer and HP’s DirectModel, have some support for these features and should be further investigated. Determining these areas early in the visualization process can also avoid unnecessary computations or data access. For instance, if we are trying to examine an iso-surface of an expensive derived variable, then calculating this variable and contour surface for regions the user never examines can be wasteful.

Impact

Unknown.

Previous Work

Unknown.

See Also: Hierarchical and Multi-resolution Techniques, Data Servers

Virtual Reality

Virtual Reality has received substantial press and hype over the past several years. As a result, the term virtual reality has been applied to almost any interactive technique or multi-modality user interface. Our definition of virtual reality requires three key components: interactivity, navigation and immersion. Interactivity is fundamental to providing the illusion of immersion. As such, much higher interactive rates are required for virtual reality than for traditional scientific visualization. The feeling of immersion also includes the use of various display devices, user input devices and user navigation devices. Traditional VR displays, such as head-mounted devices and the CAVE, incorporate stereo graphics and a full field of view. The ability to smoothly navigate through this space is also key to providing a sense of reality. This navigation is still somewhat cumbersome for gross movements with most current systems. Many new input devices are being developed to address this. Devices are also being developed to provide haptic (touch) feedback, as well as directional sound. These user interfaces are important for ASCI as well, as is the navigation. The utility of the multi-modalities is uncertain. Since most of the technologies addressed in this report aim at providing better interaction, all of these can be viewed as a foundation for getting to a virtual reality system. The trade-off between accuracy and interactivity still needs to be resolved in this setting, and virtual reality places more severe constraints on the acceptable frame update rates.

Impact

As many of the technologies note, interactivity is very important. Whether the immersive and navigational aspects of virtual reality provide as much benefit is still unknown.

Previous Work

Unknown.

See Also:

Visual Metaphors

Several techniques exist to visualize three-dimensional scalar fields. Fewer techniques exist for the representation of vector fields, tensor fields and multi-variate data fields. Research has lately developed new techniques for scalar visualization, such as volume rendering, and techniques for vector field visualization. More research is needed into techniques for representing scalar and vector fields on irregular grids, representing tensor fields, and representing multi-variate data.

Impact

Unknown.

Previous Work

Unknown.

See Also: Comparative Techniques, Image-based Rendering, Time-varying Analysis, Uncertainty and Error Visualization

Attendees:

James Ahrens
Advanced Computing Laboratory, MS B287
Los Alamos National Laboratory
Los Alamos, NM 87545
Phone: (505) 667-5797
Fax: (505) 665-4939
E-mail: ahrens@acl.lanl.gov

Michael Cox
NASA Ames Research Center
MS T27A-2
Moffett Field, CA 94035-1000
Phone: (650) 604-27225
Fax: (650) 966-8669
E-mail: mbc@nas.nasa.gov

Roger Crawfis
Assistant Professor
Department of Computer and Information Science
The Ohio State University
395 Dreese Lab
2015 Neil Avenue
Columbus, Ohio 43210-1277
Phone: (614) 292-2566
Fax: (614) 292-2911
E-mail: crawfis@cis.ohio-state.edu

Mark Duchaineau
Lawrence Livermore National Laboratory
Livermore Computing
P.O. Box 808, L-561
Livermore, CA 94550
Phone: (510) 423-1320
Fax: (510) 423-8704
E-mail: duchaineau1@llnl.gov

Chuck Hansen
Research Associate Professor
Computer Science Department
University of Utah
3190 Merrill Engineering Building
3337 MEB
Salt Lake City, Utah 84112
Phone: (801) 581-3154
Fax: (801) 581-5843
E-mail: hansen@cs.utah.edu

Carol Hunter
Assistant Deputy Department Head
Scientific Computing and Communications Department
Lawrence Livermore National Laboratory
P. O. Box 808, L-60
Livermore, CA 94550
Phone: (510) 422-1657
Fax: (510) 423-8715
E-mail: chunter@llnl.gov

Nelson Max
University of California, Davis
Livermore Department of Applied Science
Lawrence Livermore National Laboratory
P.O. Box 808, L-307
Livermore, CA 94550
Phone: (510) 422-4074
Fax: (510) 423-4139
E-mail: max2@llnl.gov

Bernd Hamann
Associate Professor
Co-Director, Center for Image Processing and Integrated Computing
University of California, Davis
3055 Engineering II
Davis, CA
Phone: (916) 754-9157
Fax: (916)
E-mail: hamann@cs.ucdavis.edu

Philip Dale Heermann
Department 09215
Sandia National Laboratory
P.O. Box 5800, MS 0318
Albuquerque, NM 87185-0318
Phone: (505) 844-1527
Fax: (505) 844-2415
E-mail: pdheerm@sandia.gov

Kathryn R. Hughes
Sandia National Laboratory
Department 08411
P.O. Box 969, MS 9101
Livermore, CA 94550-9101
Phone: (510) 294-2902
Fax:(510) 294-1539
E-mail: krhughe@sandia.gov

Mark Miller
Scientific Computing Applications
Lawrence Livermore National Laboratory

P.O. Box 808, L-170
Livermore, CA 94550
Phone: (510) 423-5901
Fax: (614) 292-2911
E-mail: Miller86@llnl.gov

Dan Schikore
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
P.O. Box 808, L-561
Livermore, CA 94550
Phone: (510) 424-5799
Fax: (510) 423-2993
E-mail: schikore1@llnl.gov

Rebecca Springmeyer
Scientific Computing &
Communications Department
Lawrence Livermore National Laboratory
P.O. Box 808, L-561
Livermore, CA 94550
Phone: (510) 423-0794
Fax: (510 423-8704
E-mail: springme@llnl.gov

Peter L. Williams, Ph.D
HO-C10
IBM Corporation
P.O. Box 392
30 Saw Mill River Road
Hawthorne, NY 10532-1507
Phone: (914) 784-7717
Fax: (914) 784-7667
E-mail: williams@cs.vassar.edu

Samuel P. Uselton
MRJ Technology Solutions
MS T27A-1
NASA Ames Research Center
Moffett Field, CA 94035
Phone: (650) 604-3985
Fax: (650) 966-8669
E-mail: uselton@nas.nasa.gov

Presentors:

Randy Christensen
Defense Nuclear Technology
Lawrence Livermore National Laboratory
P.O. Box 808, L-66
Livermore, CA 94550
Phone: (510) 423-3054
Fax: (510) 423-4820
E-mail: christen@k2.llnl.gov

Charles McMillan
Defense Nuclear Technology
Lawrence Livermore National Laboratory
P.O. Box 808, L-170
Livermore, CA 94550
Phone: (510) 423-4994
Fax: (510) 422-3389
E-mail: mcmillan@viper.llnl.gov

Douglas S. (Doug) Miller
Defense Nuclear Technology
Lawrence Livermore National Laboratory
P.O. Box 808, L-016
Livermore, CA (94550
Phone: (510) 424-4738
Fax:
E-mail: doug@guppy.llnl.gov

Kris Winer
Defense Nuclear Technology
Lawrence Livermore National Laboratory
P.O. Box 808, L-170
Livermore, CA 94550
Phone: 510) 422-2419
Fax: (510) 422-3389
E-mail: winer1@viper.llnl.gov