DOE Early Career Principal Investigator Award

An End-to-End Processing Pipeline for Large Scale Time-Varying Data Visualization

DOE Award No. DE-FG02-03ER25572

Han-Wei Shen
Department of Computer Science and Engineering
The Ohio State University
hwshen@cse.ohio-state.edu
(614) 292-0060

I. DOE Collaborations

During the past project years, we have worked closely with DOE scientists. Specifically, we have been working with the following key personnel:

(1) Ross Toedte at Oak Ridge National Laboratory on the Tera-scale Supernova Initiative (TSI) project: Our role on this project is to develop volume rendering algorithms for adaptive mesh refinement (AMR) data. The PI is listed as an official collaborator for the TSI project. One of PI's graduate students, Chaoli Wang, have also worked at Oak Ridge as an intern during summer 2004 and 2005.

(2) Dr. Jinzhu Gao and Dr. Jim Arthur Kohl at Oak Ridge National Laboratory: We collaborate with Dr. Gao and Dr. Kohl on parallel volume rendering for time-varying data and time-dependent visibility culling. A joint paper on time-varying visibility was presented and well received in IEEE Visualization 2004 conference last October.

(3) Dr. Jim Ahrens at Los Alamos National Laboratory: The goal is to develop a parallel streaming architecture for time-varying data to support the needs of Los Alamos Scientists. One graduate student of the PI, Jonathan Woodring, had worked as a summer intern in Los Alamo in 2005.

II. Project Goal

Although intensive research has been undertaken to optimize the performance of visualizing very large data sets, most of the existing methods have not targeted time-varying data. In this project, we will perform a comprehensive study of an end-to-end solution to facilitate efficient processing of large-scale time-varying data. More specifically, in this research, we will study the design of a high performance processing pipeline for time-varying data visualization with a special focus on the following research components (1) Lossless spatio-temporal data encoding and indexing schemes allowing for interactive visualization of time-varying data at arbitrary spatial and temporal scales (2) Run time data reduction, including efficient visibility culling of time-varying isosurfaces, and a compact indexing scheme to identify data blocks that contain time-varying isosurfaces. (3)Accelerated visualization algorithms and high performance parallel algorithms, which utilize temporal coherence and perform dynamic selection of spatial and temporal level of detail. (4)Time-varying feature extraction and enhancement, including high dimensional data integration, projection, and time-varying volume transfer functions.

III. Major Research Activities and Impacts

A. Time-Varying Data Management Framework for Multi-resolution Time-Varying Data Encoding and Rendering We construct a general yet efficient data management framework for visualizing large scale time-varying data. We are primarily focused on multi-resolution data management and encoding of time-varying data, with a focus on deploying such a framework on parallel cluster computers. We have designed a wavelet based Time-Space Partitioning Tree (WTSP Tree), and we have augmented WTSP trees with a load-balanced data distribution scheme to guarantee that for any arbitrary spatial and temporal resolutions selected by the user, a good load distribution among the parallel processors can be achieved so a near-optimal parallel speedup can be gained.

B. Asynchronous Rendering of Time-Varying Volumes with Workload Based Data Shuffling Given that sometimes it is difficult to balance the rendering workload within a single time step, we are developing algorithms that can perform more effective load-balancing over a period of time. We developed a workload based data shuffling scheme with asynchronous rendering, where processors do not wait for all the other processors to finish and start compositing, but rather continue rendering frames under a shuffled distribution. This distribution estimates the workload and cost of data partitions based on the derivatives of the bricks, and assigns the partitions to processors based on total costs, so that high costs go where low costs have been previously assigned; and vice versa.

C. Time-Varying Feature Tracking Feature tracking is a useful method for visualizing and analyzing time-varying scalar fields. It allows scientists to focus on regions of interest and track their evolution and interaction over time. To allow the user to freely explore the data set, features must be tracked in an efficient manner. In the last project period, we have developed an efficient time-varying isosurface tracking algorithm. Unlike the previous algorithms which compute the corresponding isosurface components in the adjacent time steps by performing expensive computation at run time, our algorithm can rapidly identify corresponding isosurfaces by performing simple table lookup operations. This table, called the correspondence lookup table, can be computed at a preprocessing stage. The idea behind our approach is that the correspondence relationship can only change at critical isovalues in R3 or R4 and remains unchanged between adjacent pairs of critical isovalues. With our algorithm, isosurfaces can be tracked in an efficient manner with minimal overhead.