

UC Davis

IDAV Publications

Title

Real-Time Out-of-Core Visualization of Particle Traces

Permalink

<https://escholarship.org/uc/item/1483p68n>

Authors

Bruckschen, Ralph W.
Kuester, Falko
Hamann, Bernd
[et al.](#)

Publication Date

2001

Peer reviewed

Real-time Out-of-Core Visualization of Particle Traces

Ralph Bruckschen*, Falko Kuester†, Bernd Hamann*, and Kenneth I. Joy*

*Visualization and Graphics Research Group
Center for Image Processing and Integrated Computing
Department of Computer Science
University of California, Davis, CA 95616-8562

† Visualization and Interactive Systems Group
Department of Electrical and Computer Engineering
University of California, Irvine, CA 92697-2625

Abstract

Visualization of particle traces provides intuitive and efficient means for the exploration and analysis of complex vector fields. This paper presents a method suitable for the real-time visualization of arbitrarily large time-varying vector fields in virtual environments. We describe an out-of-core scheme in which two distinct pre-processing and rendering components enable real-time data streaming and visualization. The presented approach yields low-latency application start-up times and small memory footprints. The described system was used to implement a “volumetric fog lance,” which can emit up to 60000 particles into a flow field while maintaining an interactive frame rate of 60 frames per second. All algorithms were specifically designed to support commodity hardware. The proof-of-concept system is running on a low-cost Linux workstation equipped with a 120GB E-IDE RAID (Redundant Array of Inexpensive Disk) system.

Keywords: Particle Tracing, Scientific Visualization, Computational Fluid Dynamics, Out-of-Core Visualization, Virtual Reality.

1 Introduction

Numerical simulations are becoming more and more powerful, simulating physical phenomena at ever increasing resolution. No longer is it possible to visualize the resulting massive data sets using an exclusive in-core approach. The goal is to provide a scientist with intuitive and interactive tools for the exploration and analysis of large time-varying data sets. The challenge is to visualize the data generated by high-precision simulation runs without loss of accuracy. Unfortunately, available computational resources on the visualization side frequently vary drastically from those on the simulation side. This implies that the amount of data has to be reduced to allow for reasonable visualization times. Using standard approaches, an increase in rendering speed generally comes at the cost of a loss in data precision. In particular when techniques such as a sub-sampling are used, undesired artifacts can be introduced or potentially important features removed. The visualization problems associated with very large time-varying data sets in computational fluid dynamics (CFD) are frequently addressed by creating image sequences or movies in batch processes. These approaches usually require basic knowledge of the result, frequently derived from interaction with lower resolution datasets. Some improvements were made with 3D interactive movies, but interactive visualization in full resolution is still expensive.

*E-mail: {rwb, hamann, joy}@cipic.ucdavis.edu

†E-mail: fkuester@ieee.org




Figure 1: Real-time particle visualization.

The goal of the presented work is to enable visualization of complex scientific data sets while using minimal CPU and memory footprints, enabling the system to run on low-cost commodity hardware as well as massively parallel systems. The discussed out-of-core approach is specialized for real-time visualization of large scale time-varying data sets at the available highest level of resolution and suitable for virtual wind tunnel scenarios. The described system was used to implement a volumetric particle injector, which can emit up to 60000 particles into the flow field while maintaining an interactive frame rate of 60 frames per second. This performance is achieved by pre-calculating a dense uniform rectilinear grid of particle injectors¹ (N^3 trajectory seeds) and storing the particles on all resulting trajectories out-of-core in a scheme optimized for selection of sub-grids.

For the initial test phase, data sets were generated for a classical problem: simulating the flow of a fluid around a spherical object. In our case, a Reynolds number² of about 2000 was selected to simulate the flow field around a sphere pushed slowly through water. A rectilinear grid of $128*128*1024$ was used for the computation of about 250 time steps. (The used finite-difference Navier-Stokes solver is courtesy of Wolfgang Borchers, University of Erlangen, Germany.)

2 Related Work

Visualization of time-varying fluid flows using large amounts of particles was done, for example, by Lane [1]. Based

¹A particle injector releases a constant stream of particles into a flow. For a numeric simulation, this is discretized by adding a particle for each time step of the simulation.

²The Reynolds number is proportional to the ratio of inertial force and viscous force. It is used in momentum, heat, and mass transfer computations to account for dynamic similarity.

on 3D movies, the possibility of rendering large amounts of particles in real time was demonstrated by Meiselbach [2]. Out-of-core particle tracing algorithms have been developed by Ueng et al. [3] for the most common forms of meshes found in CFD. A more detailed description of the problem definition and related issues can be found in [4]. A server-based approach supporting interactive frame rates was demonstrated by Cox and Ellsworth [5] for relatively small particle systems. The interactive visualization of flow data sets in virtual reality (VR) environments was studied by Jern and Earnshaw [6], and for time-dependent data sets by Kenwright and Lane [7] using in-core algorithms.

3 Visualization Concepts

The classical visualization approach for interactive particle tracing is to read the data needed to calculate the particle traces from a storage system and to calculate particle trajectories procedurally and in-core. The velocity of the particles is then interpolated and their new positions are calculated. Using this method, visualization with artificial fog is impractical, as very large massive parallel systems are needed. This is due to the amounts of data needed for the integration. The main limiting factor is data throughput from the storage system to the server.

In a time-dependent CFD simulation on an unstructured tetrahedral grid, the minimal data required to compute a particle path are two tetrahedra with velocity vectors defined at their corners. This means that for every particle two sets of four coordinates with four velocity vectors are needed. When using single-precision floating-point variables, 256 bytes will have to be transferred per particle. Assuming a static mesh, this amount can be cut in half when the mesh is cached in main memory. The other limiting factor is the amount of searches required on the storage system. As “seeks” generally require more time than the actual read operations, usually even bigger data blocks need to be processed for each particle.

As all these methods strive to achieve complete freedom of interaction for user specifiable input, they are optimized for random access of data on the storage system. For actually rendered particles, the amount of data needed is much smaller. Every particle can be described by its coordinates. Using single-precision floating-point variables, twelve bytes are needed. By transforming the coordinates to 16-bit integer values, this amount can be reduced to six bytes. Sometimes, particles are color-coded by an underlying scalar field value such as pressure, or absolute speed. For this purpose, an additional 16-bit value can be added to align the data to 64 bits (8 bytes). Thus, particles need only $\frac{1}{16}$ of data throughput compared to the data needed to actually generate them.

As complete freedom of interaction is rarely needed, the method described here limits interaction by pre-calculating particle traces in a raster. Instead of calculating particle traces on the fly procedurally, and traversing the complete visualization pipeline in real time, our approach calculates particle traces in bulk large enough to allow interaction by selection. It is possible to store the particles in a scheme so that they can be interactively selected in real time during later data exploration. The drawback is that a new raster needs to be selected and the particles re-calculated to explore different parts of the flow. This turn-around time highly depends on the size of the data, its data's structure, and the used computational hardware. As particle tracing can be parallelized efficiently, either a simulation mainframe or a

render farm could be used. The simulation we have used for testing is a sequential finite-difference method applied to a rectangular grid including a particle tracer. After each time step, new particles are injected from a raster of entry points. All particles in the integration area are traced and their coordinates, time of entry, and number of entry points are stored. The integration scheme is a second-order von Heun method using the two time steps that are stored in main memory during the simulation. The second-order scheme is sufficient, as the simulation itself is only of first order. The time stepping of the particle integration is the same as the stepping of the simulation. This keeps the particles synchronized with the simulation speed.




Figure 2: System Architecture

4 Algorithms

The particle visualization system is split into two main components: the particle tracer & encoder and the renderer. We were able to instrument the existing particle tracer allowing us to convert the particle data into the proposed scheme.

4.1 Particle Tracer & Encoder

Particle injectors are located in a uniform rectilinear mesh that discretizes a volume in a user-specified region. The volume can be intersected with the model located in the flow field. When this is done, the particle injector will be disabled for the impacted region. For every injector, the particles are stored in a list, sorted by time of entry. Every time step, a particle is added to the list with the coordinates of the injector, and new particle positions are calculated. All particles that leave the integration space will be deleted. When the maximum particle threshold is reached, particles are deleted based on their age, i.e., in FIFO order. This approach simulates nicely the dissolving of paint molecules in a liquid.

During the encoding step, the maximum amount of particles per injector is calculated, and all active injectors are sorted by their Morton code³. According to the resulting order, the particle lists for each injector are then traversed.

³The Morton order defines a linear numbering of the leaves of a complete quadtree (2D case) or octree (3D case) [9]. A Morton order optimizes the amount of searches needed to select a rectangular subset of a uniform rectilinear grid. In the 3D case, the worst case requires eight searches to access a rectangular subset, see Figure 3. The maximum size of the data to be read is eight times the size of a rectangular subset.

particles is reached at time step 130, when the broken-down vortex is carried out of the integration area.




Figure 4: Vortex structure.




Figure 5: Vortex break down.




Figure 6: Double helix structure.

5.2 Encoder

The encoder was implemented in C and utilizes the local RAID on the Linux system. Every time step is read into main memory, sorted, re-sampled, padded, and written into consecutive files. After every time step, the header file is also generated and stored. The file names are directly correlated with the current time step. Since the converter parses ASCII data piped to it from a de-compression process, each time step currently requires about five minutes to complete. The padding penalty for missing particles is up to five megabytes per time step, resulting in a maximum file size of 28MB for the described problem. This overhead, however, enables significantly faster look-up operations on block-aligned data. The alternative would be to store block sizes per time step and injector in main memory. Using the 76^3 grid with 250 time steps, this would result in an overhead of approximately 1.4 gigabytes to be stored in main memory, which conflicts with the original goal to work with the smallest possible

CPU and main memory foot prints. The resulting data set for the described test case is 60GB.

5.3 Renderer

The renderer is implemented as a plug-in component for VirtualExplorer, a VR toolkit developed at CIPIC [8]. VirtualExplorer (VE) is a customizable plug-in-based VR framework for immersive scientific visualization, data exploration and geometric modeling. The framework is layered on top of a run-time plug-in system and re-configurable virtual user interface and provides a variety of plug-in components. The system enables access to scene-graph-based APIs, including OpenGL Performer and OpenInventor, direct OpenGL support for visualization of time-critical data as well as collision and generic device managers. Plugins can be loaded, disabled, enabled, or unloaded at any time, triggered either through pre-defined events or through an external Python-based interface. The framework is currently being extended with a variety of application areas in mind, but its main emphasis is on user-guided data exploration and high-precision engineering design.

The CFD visualization toolkit provides two distinct plug-ins to enable out-of-core particle visualization. An optimized data loader and data streaming component and a designated OpenGL-based real-time rendering engine.

The required user interface and interaction metaphors are provided through the VE framework. For example, the virtual fog lance can be either controlled through a spatially tracked input device or its appropriate keyboard- or software-based simulator.

At start-up time, the plugins are launched as independent processes. The data acquisition process then acquires the data from the RAID system and stores it in shared memory, together with the needed pointers. The specific data sets are determined by the current spatial coordinates of the cursor that is provided through the VE framework and constantly dispatched to the plugins.

In the rendering phase, the current particle coordinates are retrieved from shared memory and processed. Currently, particles are rendered as simple points with user adjustable color. The blend function of the 3D hardware is set to add up the RGB values per rendered pixel, thus achieving the desired fog effect. Particles that are in alignment with the viewing direction are brighter. After all particles are rendered, a scene graph is traversed to visualize the cursor position, the model and the integration area.

5.4 Hardware Setup

All algorithms were developed with a wide range of system configurations in mind. However, the emphasis was on providing real-time visualization and interaction capabilities for low-cost commodity hardware. The current proof-of-concept system runs under RedHat Linux 6.2 on a PentiumIII (866Mhz) system with 512MB of main memory and a GeForce2 GTS graphics board (64MB). The data sets are stored on a 120GB E-IDE RAID system using four low-cost IBM DeskStar ATA 100 disks and an AMI Hyperdisk soft-raid controller. At this point, the available Linux drivers limit the RAID to ATA 66. Considering the total price of less than US\$ 2500 for this setup, the system delivers an astonishing price-performance ratio.

6 Results

A uniform rectilinear particle injector grid of 76^3 has shown to be sufficient to visualize the key features of the turbulent flow. However, further tests with different data sets are required to determine if an “optimal” injector size can be defined for arbitrary data.

Initial runs indicate that the visualization algorithms are bound by the transfer rate from the side of the RAID. The data transfer rate currently translates into a stable update rate of 10 frames per second, with each frame representing a unique time step. The average CPU load on the described system is roughly 20% while utilizing 128MB of RAM.

The available data update rate of 10Hz is sufficiently high to support smooth visualization. While the update rate is bound by the speed of the RAID system, the threaded rendering process runs at a frame rate of more than 75 frames per second. This performance makes the system suitable for real-time interactive stereoscopic rendering in virtual environments.

The under-utilized CPU and memory resources allow for a variety of additional improvements in rendering techniques and interaction. However, the turn-around time of the simulator using the current sequential implementation needs to be improved by a factor of at least 10.

6.1 Performance

At the moment, the performance of the pre-processor seems to be poor in comparison with the performance of the renderer. As particle tracing can be efficiently parallelized, this is less concern.

6.1.1 Pre-processing

Even though particle tracing for rectangular meshes leads to very fast algorithms, the old sequential technique integrated in the finite-difference simulation program requires about a day to calculate all particles for the 250 time steps of the simulation. An improved parallel, out-of-core particle tracer is needed and should exploit bulk particle rendering.

6.1.2 Rendering

The theoretical minimal refresh rate of this system is given by

$$\frac{1}{8t_{seek} + t_{maxread}},$$

where t_{seek} is the average seek time of the RAID system and $t_{maxread}$ is the time needed to read in the largest block of particles in seconds. The value of $t_{maxread}$ can be estimated from the raw throughput of the RAID system.

In our case, the assumed value of t_{seek} would be 10ms. With a cursor size of 8^3 injectors and an average of 130 active particles per injector, the largest data block size is 520KB assuming that eight bytes are used to represent each particle. With a data transfer rate of 60MB per second, the RAID system can deliver 118 blocks per second, resulting in a value of $t_{maxread}$ of 8.47ms.

Thus, a theoretical minimal update rate of 11 frames per second is possible. The achieved frame rate is currently on the lower theoretical limit. An open question is whether the 60MB per second peak transfer rate can be achieved continuously or whether the used RAID system is suffering from fragmentation.

7 Conclusions

We have presented a method for the interactive visualization of particle traces for large time-dependent CFD data sets and demonstrated that this method can be performed on commodity hardware. Our out-of-core visualization technique provides interactive frame rates and scales for growing data set sizes.

8 Future Work

A client-server based implementation is being currently investigated, which will enable access to a wide range of available VR visualization hardware, such as the ImmersiveWorkbench. The achieved data rates suggest that a coupling of the current system using the common 100Mbit Ethernet is feasible.

System tests on a wide variety of already existing data sets are in progress. These tests will determine how efficient the data conversion tools are and how the system performs with tera-byte data sets streamed over a high-speed network.

Experiments on how this method scales on larger and faster RAID systems need to be conducted and a wide range of possible optimization strategies such as data set partitioning and hard disk de-fragmentation explored. Different algorithms for the creation of photo-realistic particle clouds will be added in the future. We will be testing our system on a multi-processor machine with a fiber channel RAID as well.

Acknowledgments


This work was supported by the National Science Foundation under contracts ACI 9624034 (CAREER Award), through the Large Scientific and Software Data Set Visualization (LSSDSV) program under contract ACI 9982251, and through the National Partnership for Advanced Computational Infrastructure (NPACI); the Office of Naval Research under contract N00014-97-1-0222; the Army Research Office under contract ARO 36598-MA-RIP; the NASA Ames Research Center through an NRA award under contract NAG2-1216; the Lawrence Livermore National Laboratory under ASCI ASAP Level-2 Memorandum Agreement B347878 and under Memorandum Agreement B503159; the Lawrence Berkeley National Laboratory; the Los Alamos National Laboratory; and the North Atlantic Treaty Organization (NATO) under contract CRG.971628. We also acknowledge the support of ALSTOM Schilling Robotics and SGI. We thank the members of the Visualization and Graphics Research Group at the Center for Image Processing and Integrated Computing (CIPIC) at the University of California, Davis. We also thank Wolfgang Borchers from the University of Erlangen, Germany, for providing us with his CFD data set.

References


- [1] D. A. Lane, “UFAT - A particle tracer for time-dependent flow fields”, in *Proceedings of the Conference on Visualization*, R. Daniel Bergeron and Arie E. Kaufman, Eds., Los Alamitos, California, USA, Oct. 1994, pp. 257–264, IEEE Computer Society Press.
- [2] C. Meiselbach and R. Bruckschen, “Interactive visualization: On the way to a virtual wind tunnel”, in

Flow Simulation with High-Performance Computers, II, Ernst Heinrich Hirschel, Ed., vol. 52 of *Notes on Numerical Fluid Mechanics*, pp. 203–208. F. Vieweg and Sohn, Wiesbaden, Germany, 1996.


- [3] Shyh-Kuang Ueng, C. Sikorski, and Kwan-Liu Ma, “Out-of-core streamline visualization on large unstructured meshes”, *IEEE Transactions on Visualization and Computer Graphics*, vol. 3, no. 4, pp. 370–380, Oct./Dec. 1997.
- [4] Frits H. Post and Theo van Walsum, “Fluid flow visualization”, Technical Report DUT-TWI-92-69, Delft University of Technology, Department of Technical Mathematics and Informatics, Delft, The Netherlands, 1992.
- [5] Michael B. Cox and David Ellsworth, “Application-controlled demand paging for Out-of-Core visualization”, in *IEEE Visualization '97*, Roni Yagel and Hans Hagen, Eds., Los Alamitos, California, USA, Nov. 1997, IEEE Computer Society Press, pp. 235–244.
- [6] Mikael Jern and Rae A. Earnshaw, “Interactive real-time visualization system using a virtual reality paradigm”, in *Visualization in Scientific Computing*, M. Göbel, H. Müller, and B. Urban, Eds., pp. 174–189. Springer-Verlag, Vienna, Austria, May 1994.
- [7] David N. Kenwright and David A. Lane, “Interactive Time-Dependent Particle Tracing Using Tetrahedral Decomposition”, *IEEE Transactions on Visualization and Computer Graphics*, vol. 2, no. 2, pp. 120–129, June 1996.
- [8] Falko Kuester, Bernd Hamann, and Kenneth I. Joy, “Virtualexplorer: A plugin-based virtual reality framework”, in *Proceedings of SPIE*, R.F. Erbacher, P.C. Chen, M. Groehn, J.C. Roberts, and C.M. Wittenbrink, Eds., San Jose, California, USA, 2001, SPIE - The International Society of Optical Engineering.
- [9] H. J. Samet, *Design and analysis of Spatial Data Structures: Quadtrees, Octrees, and other Hierarchical Methods*, Addison-Wesley, Redding, MA, 1989.
- [10] T. J. Chung, *Finite Element Analysis in Fluid Dynamics*, McGraw Hill, 1987.
- [11] Robert Haines and Dave Darmofal, “Visualization in computational fluid dynamics: A case study”, in *Proceedings of Visualization '91*, Los Alamitos, California, USA, 1991, pp. 392–397, IEEE Computer Society Press.
- [12] W. Thompkins and R. Haines, “A minicomputer/array processor/memory system for large-scale fluid dynamic calculations”, in *Impact of New Computing Systems on Computational Mechanics*, A. Noor, Ed., pp. 117–126. The American Society of Mechanical Engineers, 1983.



Color Plate 1: The vortex structure at the beginning of the simulation. As the sphere is suddenly moved through the liquid, a fog ring forms.



Color Plate 2: The vortex breaks down after about 50 timesteps.



Color Plate 3: The vortex is replaced by a helix like structure.

