SAND2010-2137P

In-Situ Visualization with the ParaView Coprocessing Library

Abstract

In-situ visualization is a term used to describe running a solver in tandem with visualization. By coupling these together in the same execution we can utilize the high performance computing for post processing, and we can circumvent the bottlenecks associated with storing and retrieving data in disk storage. To simplify the integration of visualization and post processing into computational code, we present the coprocessing library provided as part of the ParaView framework. ParaView is a powerful open-source turnkey application for analyzing and visualizing large data sets in parallel. The coprocessing library provides a programmatic interface to the ParaView framework that is designed to simplify integration with existing codes. This tutorial presents the coprocessing library and the fundamentals of how to use it. Attendees will learn the structure of the coprocessing API and how to bind it to C, C++, FORTRAN, and Python. Attendees will also receive instructions on customizing the ParaView build to a particular code, minimizing the memory footprint, and simplifying the specification of analysis tasks via integration with the ParaView application.

Detailed Description

General description and tutorial goals

We are proposing a half-day tutorial on adding post processing and visualization to existing codes. The tutorial presents the interface presented by the ParaView coprocessing library and the visualization and analysis capabilities it provides. Attendees learn how to integrate the coprocessing library into their existing codes and to configure and apply the analysis.

Targeted audience, content level, and relevance to SC attendees

We expect the content break down to be: 10% beginner, 40% intermediate, 50% advanced.

This tutorial has an appeal for many levels and types of attendees, but is primarily focused on those working with large-scale simulation code. Attendees currently programming simulations or other large-scale codes will directly benefit from the expansive capabilities of the ParaView coprocessing library, and attendees familiar with visualization, graphics, or VTK will learn how to programmatically interface to these components in large-scale applications.

Beginners benefit by getting an overview of the capabilities of the ParaView framework and an introduction to applying it. This understanding helps them choose and apply their visualization tools. The tutorial also provides the necessary information for using the coprocessing tools already integrated into other codes.

Intermediate and advanced users are given the instruction necessary to leverage the ParaView framework within their codes. We give an overview of the coprocessing library's API and detailed instruction on how to use it. We also apply instruction on how to configure and apply the analysis within the ParaView framework, and how to create a workflow that provides analysis both running *in situ* with the simulation and interactively at a user's convenience.

The majority of the Supercomputing attendees are from universities or government labs. In this environment it is important to be able share tools and applications. Since this suite of tools is open source, there are no barriers to collaboration between diverse organizations. The fact that there are no license fees for these applications is also important. Some university research groups may not be able to easily purchase expensive licenses for proprietary visualization applications.

Also, ParaView is a world leader in high performance visualization on distributed clusters. Expensive shared memory computers have succumbed to the economically superior distributed cluster computer. Many Supercomputing attendees are researching the application of distributed clusters toward high performance computation and visualization. ParaView has been designed from inception to run well on distributed computing platforms.

Visualization and analysis are important tools for debugging simulations, sharing results and marketing supercomputing research. Knowledge of these flexible and configurable

tools will benefit a wide audience of conference attendees (not only ones specializing in visualization).

Audience prerequisites

Much of the discussion in this tutorial involves using the programming interface to the coprocessing library. Attendees will get the most out of the tutorial if they have enough programming experience to follow along with the discussions. The tutorial provides examples in C++, FORTRAN, and Python. For our target audience, familiarity with at least one of these languages is prevailing.

Ensuring cohesive content

The presenting institutions, Sandia National Laboratories and Kitware, Inc., have worked closely together in the development of much of the software presented in this tutorial, and will continue to work closely together in building the tutorial content.

Tutorial updates for SC

We have given many introductory tutorials to ParaView in different forums, including several past Supercomputing conferences. However, this tutorial deviates significantly from past tutorials. Whereas previous tutorials focused primarily on the end user application (the GUI) of ParaView, this tutorial focuses on a *programmatic* interface into ParaView. The intent is to provide a service to those developing other large-scale parallel applications, and Supercomputing is an excellent opportunity for that.

This tutorial is more advanced in the sense that, because the interface is programmatic, it requires a programming skill set that was not necessary before. But the tutorial still provides beginning instruction to the visualization algorithms and systems, so in that regard it is still introductory nonetheless.

Acknowledgements

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Description of Demonstration

The information in this tutorial is communicated via slide presentations and live demonstrations, both from a laptop provided by the presenter.

Required Hardware

Standard A/V equipment is sufficient for this tutorial.

Outline

1. Introduction (30 minutes).
2. Coprocessing API (1 hr 30 minutes).
    a. Initialize/Finalize.
    b. Establishing a pipeline.
        i. Python scripting.
        ii. Hard coding.
    c. Specifying/querying input data.
    d. VTK data structures.
    e. Linking to FORTRAN.
    f. The coprocessing ParaView plugin.
3. Compiling (30 minutes).
    a. Cross compiling.
    b. Compiling without shared libraries.
    c. Compiling with and without Python.
    d. Specifying filters to minimize footprint.
4. Coprocessing as an I/O Service (30 minutes).
    a. Overview.
    b. Available I/O interfaces.

Tutorial Notes Release

The presenters of this tutorial agree to release the tutorial notes on the SC10 USB stick.

Travel Support Request

Some of the presenters will request support for travel.

Résumé: Kenneth Moreland

Education

- Doctor of Philosophy Computer Science, University of New Mexico, Albuquerque, NM, July 2004.
- Master of Science Computer Science, University of New Mexico, Albuquerque, NM, May 2000.
- Bachelor of Science Computer Science, New Mexico Institute of Mining and Technology, Socorro, NM, May 1997.
- Bachelor of Science Electrical Engineering, New Mexico Institute of Mining and Technology, Socorro, NM, May 1997.

Work Experience

ParaView 3 Development Lead — Sandia National Laboratories, 10/2006–Present

Lead the ASC funded development effort for ParaView 3, a large-scale general visualization solution.

Scalable Visualization — Sandia National Labs, 8/2000–Present

Researched and developed parallel rendering codes and other scalable visualization algorithms targeted at performing scientific visualization on cluster computers.

Product Realization Environment — Sandia National Labs, 5/1996–8/2000

Developed and deployed a CORBA-based middleware tool for distributing and connecting scientific modeling and simulation codes.

Research Experience (Academic)

Doctoral Research — University of New Mexico, 8/01–5/04

Conceived novel mathematical analyses to the direct volume rendering light transport equation that enable real time display of highly accurate, linearly interpolated unstructured volumes.

Selected Publications

- Kwan-Liu Ma, Chaoli Wang, Hongfeng Yu, Kenneth Moreland, Jian Huang, and Rob Ross. "Next-Generation Visualization Technologies: Enabling Discoveries at Extreme Scale." *SciDAC Review*, Issue 12, pg. 12–21, Spring 2009.
- John Biddiscombe, Berk Geveci, Ken Martin, Kenneth Moreland, and David Thompson. "Time Dependent Processing in a Parallel Pipeline Architecture." *IEEE Transactions on Visualization and Computer Graphics*, Volume 13, Number 6, pg. 1376–1383, November/December, 2007. DOI=10.1109/TVCG.2007.70600.

- Kenneth Moreland, Lisa Avila, and Lee Ann Fisk. "Parallel Unstructured Volume Rendering in ParaView." In *Visualization and Data Analysis 2007, Proceedings of SPIE-IS&T Electronic Imaging*, pg 64950F-1–12. January 2007.
- Andy Cedilnik, Berk Geveci, Kenneth Moreland, James Ahrens, and Jean Farve. "Remote Large Data Visualization in the ParaView Framework." In *Eurographics Parallel Graphics and Visualization 2006*, pg. 163–170. May 2006.
- Kenneth Moreland and David Thompson. "From Cluster to Wall with VTK." In *Proceedings of IEEE 2004 Symposium on Parallel and Large-Data Visualization and Graphics*, pages 25–31. October 2003.
- Kenneth Moreland, Brian Wylie, and Constantine Pavlakos. "Sort-Last Parallel Rendering for Viewing Extremely Large Data Sets on Tile Displays." In *Proceedings of the IEEE Symposium on Parallel and Large-Data Visualization and Graphics*, pages 85–92. October 2001.
- Brian Wylie, Constantine Pavlakos, Vasily Lewis, and Kenneth Moreland. "Scalable Rendering on PC Clusters." *IEEE Computer Graphics and Applications*, volume 21, number 4, pages 62–70. July/August 2001.

Selected Presentations

- Kenneth Moreland, John Greenfield, W. Alan Scott, Utkarsh Ayachit, Berk Geveci, and David DeMarle. "Large Scale Visualization with ParaView." *Supercomputing 2008*, November 2008.
- Kenneth Moreland and John Greenfield. "Large Scale Visualization with ParaView 3". *Supercomputing 2007*. November 2007.
- "Parallel Visualization with ParaView." *Supercomputing 2005*. November 2005.
- "Large Scale Visualization with Cluster Computing." *Linux Cluster Institute Workshop*. October 2004.
- "Big Data, Big Displays, and Cluster-Driven Interactive Visualization." *Workshop on Commodity-Based Visualization Clusters*. October 2002.

Professional Affiliations

- Institute of Electrical and Electronic Engineers (IEEE)
- Association for Computing Machinery (ACM)

Résumé: John Greenfield

Education

- Doctor of Philosophy Computer Engineering, University of New Mexico, Albuquerque, NM, December 2003.
- Master of Science Computer Engineering, University of New Mexico, Albuquerque, NM, May 1992.
- Bachelor of Science Electrical Engineering, University of New Mexico, Albuquerque, NM, May 1985.

Work Experience

Visualization Support ASAP under contract to Sandia National Laboratories,

3/2004 – Present

Provide support to both developers and users of visualization support tools especially EnSight and ParaView.

Sr Research Scientist University of New Mexico High Performance Computing Center,

8/1997 – 3/2003

Researched parallel processing systems and visualization software, Supported large scale videoconferencing systems and user training.

Research Experience (Academic)

Doctoral Research University of New Mexico, 8/2001 – 10/2003

Investigated impact of directional audio on large scale videoconferencing systems.

Selected Publications

- Kenneth L. Summers, John Greenfield and Brian T. Smith. "A Survey of Parallel Program Performance Evaluation Techniques using Visualization and Virtual Reality," in Proceedings, IEEE Aerospace conference 2000.
- Robert A. Ballance, Thomas P. Caudell, John Greenfield and Rick Stevens. "The National Computational Science Alliance Access Grid: An Internet-Based Collaboration Tool Augmented by High-Performance Computing," in Proceedings of the DoD High Performance Computing Users Group Conference 2000, June 2, 2000
- S. Salvini, B.T. Smith, and J. Greenfield, "Towards mixed mode parallelism on the new model F50-Based IBM SP system," Albuquerque High Performance

Computing Center, University of New Mexico, Technical Report AHPCC98003, September 1998.

- John A. Greenfield, Richard D. Hunt, and Gregory L. Heileman. "A Parallel Load Balancing Technique for Finite Element Method Problems," University of New Mexico, Technical Report EECE95-002, February 28, 1995

Selected Presentations

- Large Scale Visualization with ParaView 3 presented at Supercomputing 2007.
- Beginning EnSight class at Sandia National Laboratories 05/01/2007
- Beginning ParaView class at Sandia National Laboratories, 02/27/2007
- "What's New in ParaView" presented at DOE Computer Graphics Forum 2005, 04/2005

Professional Affiliations

- Association for Computing Machinery (ACM)

Résumé: Berk Geveci

Education

- Doctor of Philosophy Mechanical Engineering from Lehigh University, 1999.
- Master of Science Mechanical Engineering from Lehigh University, 1996.
- Bachelor of Science Mechanical Engineering from Bogazici University, 1994.

Work Experience

ParaView Project Lead — Kitware Inc.

Lead the open source ParaView project.

Research Experience (Academic)

Post-Doctoral Fellowship — University of Pennsylvania

Worked in the area of optimal control investigating applications in the control of hydrothermal instabilities.

Doctoral Research — Lehigh University

Conducted research on subsonic and supersonic flow induced nonlinear vibrations, developing a new procedure for the solution of coupled flow and structural equations. In addition, he authored software for the study of separation in unsteady boundary layer flows and the visualization of the numerical and experimental results.

Selected Publications

- John Biddiscombe, Berk Geveci, Ken Martin, Kenneth Moreland, and David Thompson. "Time Dependent Processing in a Parallel Pipeline Architecture." *IEEE Transactions on Visualization and Computer Graphics*, Volume 13, Number 6, pg. 1376–1383, November/December, 2007. DOI=10.1109/TVCG.2007.70600.
- Andy Cedilnik, Berk Geveci, Kenneth Moreland, James Ahrens, and Jean Farve. "Remote Large Data Visualization in the ParaView Framework." In *Eurographics Parallel Graphics and Visualization 2006*, pg. 163–170. May 2006.
- James Ahrens, Berk Geveci and Charles Law. "ParaView: An End-User Tool for Large Data Visualization." In *The Visualization Handbook*. Edited by C.D. Hansen and C.R. Johnson. Elsevier. 2005.
- K.M. Martin, B. Geveci, J. Ahrens, C. Law. "Large Scale Data Visualization Using Parallel Data Streaming." *IEEE Computer Graphics & Applications*, July 2001.

Résumé: Utkarsh Ayachit

Education

- Master of Science Computer Science from University of Maryland, Baltimore County, 2004.

Work Experience

Research Staff — Kitware Inc.

Contributed to the design and implementation of ParaView.

Research Experience (Academic)

Master's Research — Lehigh University

Explored the use of level-of-detail in 2D flow visualization techniques namely streamline visualizations, vector plots and texture-based techniques.