UC Davis
IDAV Publications

Title

A Framework for Visualizing Hierarchical Computations

Permalink

https://escholarship.org/uc/item/0533s51n

Authors

Ligocki, Terry J.
Van Straalen, Brian
Shalf, John M.
et al.

Publication Date

2003

Peer reviewed

A Framework for Visualizing Hierarchical Computations

Terry J. Ligocki[1], Brian Van Straalen[2], John M. Shalf[1], Gunther H. Weber[3,4], and Bernd Hamann[1,3]

[1] Lawrence Berkeley National Laboratory, National Energy Research Scientific Computing Center, 1 Cyclotron Road, M/S 50F, Berkeley, CA 94720, USA
[2] Lawrence Berkeley National Laboratory, National Energy Research Scientific Computing Center, 1 Cyclotron Road, M/S 50A-1148, Berkeley, CA 94720, USA
[3] Center for Image Processing and Integrated Computing, Department of Computer Science, University of California, 1 Shields Avenue, Davis, CA 95616-8562, USA
[4] AG Graphische Datenverarbeitung und Computergeometrie, FB Informatik, Universitaet Kaiserslautern, Postfach 3049, D-67653 Kaiserslautern, Germany

Summary. Researchers doing scientific computations are attempting to accurately model physical phenomena. When these physical phenomena take place at a variety of different spatial and temporal scales it can be more efficient and accurate to model them at different levels of detail in an adaptive, hierarchical manner. We present a framework for visualizing adaptive, hierarchical computations – a conceptual framework and an implementation framework. Given that researchers have already defined a hierarchical structure for their data and are performing their computations using this structure, it has become important to provide a visualization tool which accurately represents this data and visualizes it directly. The tool we have designed for this purpose was built using the Visualization Toolkit, VTK, and one of its interpretive interfaces, Tcl/Tk. In addition to creating a visualization tool, we are developing extensions of visualization techniques and algorithms to hierarchical data (e.g., seamless isosurface generation).

1 Introduction

Researchers doing scientific computations are attempting to accurately model physical phenomena. When those physical phenomena take place at a variety of different scales it can computationally be more efficient and accurate to model them at different levels of detail in an adaptive manner. Two groups in the National Energy Research Scientific Computing center, NERSC [14], at the Lawrence Berkeley National Laboratory, LBNL [12], are doing just that. One group is headed by John Bell (Center for Computational Sciences and Engineering, CCSE [4]), and the other is headed by Phil Colella (Applied Numerical Algorithms Group, ANAG [2]). Both groups are doing computations using similar adaptive mesh refinement, AMR, techniques. Since the term "AMR" can mean a variety of things to researchers it should be clarified that we use it to refer exclusively to block-structured AMR data as defined in Berger and Colella [3].

Given that researchers have already defined a hierarchical structure for their data and are performing their computations using this structure, it has become important to provide a visualization tool that accurately represents and visualizes this data. This work has been a joint effort between ANAG and the LBNL/NERSC Visualization Group [17]. It includes extending and modifying visualization algorithms (e.g., isosurface computation, streamline generation) to correctly generate results while taking advantage of inherent computational efficiencies supported by the original AMR data structure. We have chosen ANAG's AMR computational library, Chombo [5], and built an extensible visualization tool, ChomboVis [6], for the data sets Chombo produces. This tool was built using the Visualization Toolkit, VTK [18], and one of its interpretive interfaces, Tcl/Tk [16]. By using VTK we have been able to use a broad foundation of existing algorithms and infrastructure while benefiting from ongoing extensions and improvements to VTK, e.g. [15]. Since VTK is an extensible, object-oriented library, we can add functionality to existing algorithms and add new algorithms relatively easily. The interpretive interface has allowed us to rapidly prototype ideas, add new functionality, benefit from the work other groups are doing with Tcl/Tk, and provide researchers using our tool with the option of directly extending the tool in an interactive fashion.

The development of this tool has proceeded in several directions. First, we have implemented and released ChomboVis, which provides researchers with many of the tools they need to view and investigate their data. This version provides 2D and 3D visualization capabilities including the ability to look at selected data directly using spreadsheets. Second, we have been developing extensions of visualization techniques and algorithms to AMR data (e.g., seamless isosurface generation). These extensions can be integrated into VTK and ChomboVis. Finally, we have been using recent extensions to VTK to handle AMR data sets in a more natural, efficient and direct manner. We believe that the new framework we have been developing provides researchers with a tool that meets their immediate needs, can and will be extended to meet future needs, and will benefit from other work being done by the visualization research community.

2   Past and Current Work

The following is a brief description of the AMR data generated by Chombo. It is intended to give the reader an idea of the structure of the data but not to be a precise or complete definition. The AMR data produced by computations using Chombo consists of a set of regular (structured) grids that are grouped by level. All grids on a given level have the same cell size or resolution. All grids on a given level are completely covered by grids on the next coarser level. The ratio of the cell size on one level to the cell size on the next finer level is always an integer. Finally, data values are cell-centered values and

the computations done on AMR grids are finite-difference approximations to partial differential equations, PDEs.

There are many approaches to extending visualization algorithms to AMR data. One of the primary difficulties is that there may be multiple data values at a given point – each value being associated with a different level in the AMR hierarchy. There are several approaches that deal with this difficulty:

1. Treat all grids (and their values) independently.
2. Use the data value(s) from the finest grid available and ignore data value(s) from the coarser grids.
3. Combine the data in some way that is physically meaningful and use the result for visualization.

Each of these approaches can be useful depending on what the user is looking for in the data sets. If users are debugging computational algorithms, the first approach allows them to look at all the data. If they are trying to understand computations and/or present results, the second or third approaches may be better.

2.1   Visualization Tool

We started developing a visualization tool by treating all grids and data independently since it was straightforward to implement and was of immediate value to ANAG researchers as they worked on the Chombo library. Also, it is of considerable value to the users of the Chombo library to have a visualization tool that can be extended. The result was ChomboVis. The first version was released in early 2000, and the second version was released in late 2000.

ChomboVis was built using VTK and its Tcl/Tk interface. VTK was chosen because it is a freely available library that includes source code and thus can be modified and extended directly. ChomboVis contains several new VTK objects – for example, an object that reads the HDF5 output of Chombo [11] and converts it into VTK data objects. Tcl/Tk was used to develop the user interface and to combine VTK objects into a working system. In addition, we envision users directly interacting with ChomboVis via its Tcl/Tk interface and creating custom extensions.

The second release of ChomboVis can be used to visualize 2D and 3D AMR data sets and provides the following capabilities (see Fig. 1 for an illustration of the tool in operation):

- Data selection by scalar variable and level
- Orthogonal data slicing and display
- Multiple isosurface/contour generation
- Spreadsheet viewing of individual grids
- Interactive selection of grids from the visualization (i.e., picking)
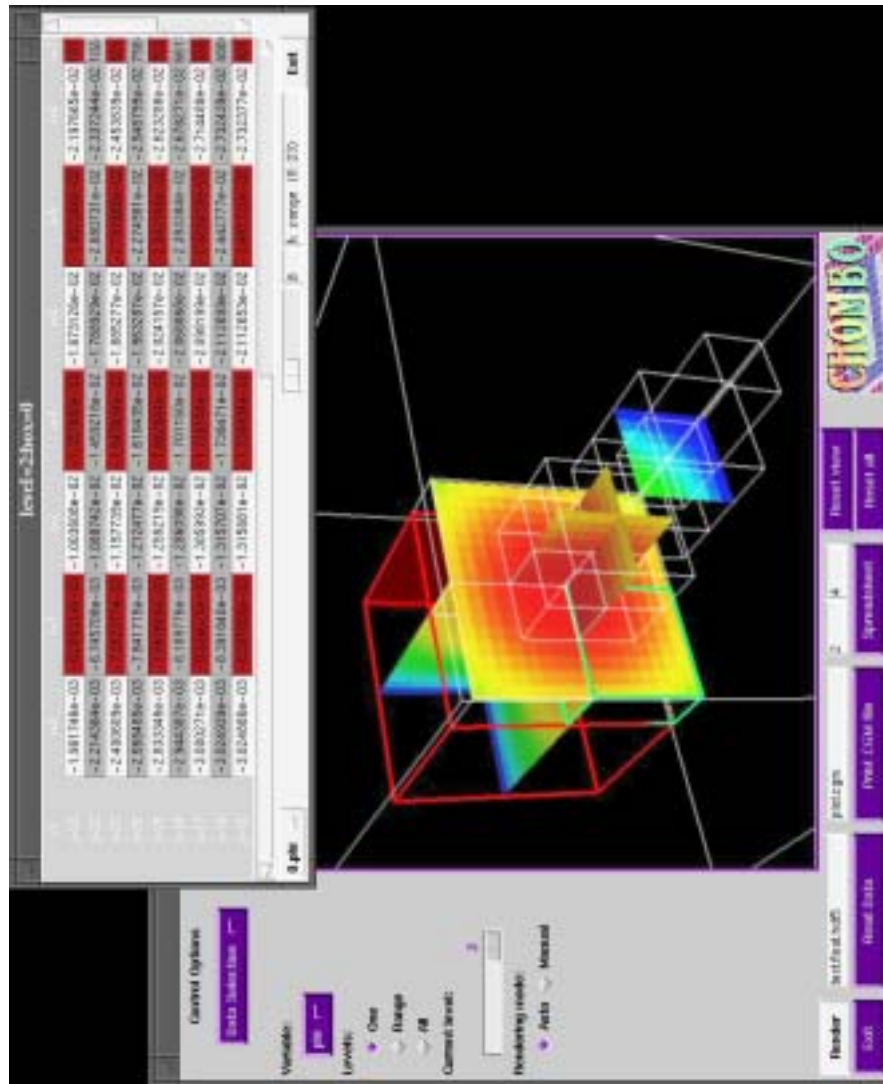- Display of grid bounds, cell size, etc.

Fig. 1. Visualization of an AMR data set using ChomboVis.

- User specified colormaps
- Output in Computer Graphics Metafile, CGM, format

These are fairly modest capabilities, and yet they required a substantial amount of development. This was due to the fact that VTK (like many widely available visualization packages) does not directly support multiple grids. Much of the work we did was to attain the goal of providing a mechanism

for sending multiple grids through a given VTK pipeline and collecting the results for rendering.

Initially we planned to use some of the recent VTK extensions to handle groups/arrays of grids [1]. This was not possible due to differences between the task the extensions were addressing (domain decomposition) and our task (handling overlapping sets of grids). We then turned to a novel, streaming, out-of-core technique base on work done by Matthew Hall [10] with VTK in order to minimize memory requirements and pipeline overhead.

2.2   Visualization Research

While the work on ChomboVis continued we began to collaborate with researchers from the Center for Image Processing and Integrated Computing, CIPIC [7], in the Department of Computer Science at the University of California, Davis. This collaboration combined nicely the work on AMR data sets done by NERSC researchers with the work done by CIPIC in hierarchical representations for and visualizations of large data sets. The goal of this collaboration was to extend visualization algorithms and techniques to AMR grids. Specifically, the following areas were and are being investigated:

1. Generation and rendering of isosurfaces from AMR data sets with no artifacts due to overlap, gaps, or cracks between grids at different levels.
2. Visualization of vector fields defined in AMR data sets.
3. Visualization of embedded boundaries, EBs, that was being developed by ANAG in conjunction with AMR data structures.
4. Interactive, immersive visualization of large AMR data sets using a variety of techniques (e.g., "seeded isosurface generation").
5. Interactive previewing (fast, with artifacts) and high-quality (slow, without artifacts) volume rendering techniques for AMR data sets.

One goal was to take advantage of the regular AMR grid structure wherever possible and only do additional work where it was necessary. This was one of the general advantages of this type of AMR representation in scientific computations. Substantial progress was made on each project. The accomplishments in each area were:

1. A technique was implemented for handling multiple, overlapping grids [8] [9]. It removes portions of grids that overlap finer grids, handles the rest using marching cubes [13], and generates stitching cells and geometry at the boundaries to create a seamless result.
2. Initial system infrastructure was developed to allow AMR vector data to be manipulated, and a technique was implemented to compute integral curves of AMR vector field data.
3. The infrastructure to represent EB data for boundary reconstruction and its visualization was implemented.

4. Interactive, seeded isosurface algorithms were developed for regular grids, curvilinear grids, and general tetrahedral meshes. Time budgets and work queues are used to achieve interactivity. This work was extended to AMR data sets.
5. Several promising techniques for interactive volume rendering were explored – including a technique that ignores much of the detailed grid structure of AMR data by viewing 3D data as point sets.

3   Future Work

There are ample opportunities for continued work both in the form of direct extensions to current work and more subtle new directions. Some of the successful visualization research will be integrated into ChomboVis, and use of ChomboVis will suggest new areas of visualization research as well.

3.1   Visualization Tool

ChomboVis is a maturing visualization tool for AMR data sets produced by Chombo. It will be used by researchers and improved to meet their needs. There are several aspects to this process. The user interface(s), documentation, and general usability will be improved. This can be done by working with users to get a better understanding of how they use ChomboVis. The performance of the tool will need to be improved for large data sets. This will require substantial modifications to VTK to handle sets of grids in a more direct fashion. There is ongoing work in this area by other users of VTK, and we plan to use this work to help us address performance issues.

We plan to integrate some of the successful AMR visualization research being done into VTK and ChomboVis to provide users with the most advanced visualization tools. Specifically, seamless isosurface generation, AMR vector visualization, and interactive volume rendering are candidates for integration.

3.2   Visualization Research

AMR visualization problems lead to many opportunities for additional research. All visualization techniques being developed could be studied in the interactive, immersive context being developed at CIPIC and NERSC. Researchers will be combining AMR and EB technologies and thus research in visualizing AMR and EB data sets will need to be combined and extended. The computing and visualizing of scalars and vectors derived from quantities originally computed is of considerable interest to researchers. Finally, applying and extending all the visualization tools and research to time-varying AMR data sets provides difficult challenges.

4   Conclusions

By working closely with researchers doing AMR computations and developing visualization tools that AMR simulation scientists use, we have developed a framework for visualizing AMR computations. This has given rise to a number of visualization research questions and problems. Work on these problems has lead to extensions of our original framework and tools. This, in turn, will lead to more fundamental visualization research.

5   Acknowledgements

References

1. Ahrens J., Law C., Schroeder W., Martin K., and Papka M. A Parallel Approach for Efficiently Visualizing Extremely Large, Time-Varying Datasets, Los Alamos National Laboratory – Tech. Report #LAUR-00-1620
2. http://seesar.lbl.gov/anag/
3. Berger M. and Colella P. (1989) Local adaptive mesh refinement for shock hydrodynamics. Journal of Computational Physics, 82:64-84, May 1989. Lawrence Livermore Laboratory Report No. UCRL-97196
4. http://seesar.lbl.gov/ccse/
5. http://seesar.lbl.gov/anag/chombo/
6. http://seesar.lbl.gov/anag/chombo/chombovis.html
7. http://graphics.cs.ucdavis.edu/
8. Weber G.H., Kreylos O., Ligocki T.J., Shalf J.M., Hagen H., Hamann B., and Joy K.I. (2001) Extraction of crack-free isosurfaces from adaptive mesh refinement data. Proceedings of the Joint EUROGRAPHICS and IEEE TCVG Symposium on Visualization, Ascona, Switzerlang, May 28-31, 2001, Springer Verlag, Wien, Austria, May 2001

9. Weber G.H., Hagen H., Hamann B., Joy K.I., Ligocki T.J., Ma K.L., and Shalf J.M. (2001) Visualization of adaptive mesh refinement data. Proceedings of the SPIE (Visual Data Exploration and Analysis VIII, San Jose, CA, USA, Jan. 22-23), Volume 4302, SPIE – The International Society for Optical Engineering, Bellingham WA, January 2001
10. http://zeus.ncsa.uiuc.edu/~mahall/
11. http://hdf.ncsa.uiuc.edu/HDF5/
12. http://www.lbl.gov/
13. Lorensen W.E. and Cline H.E. (1987) Marching Cubes: A high resolution 3D surface construction algorithm. Computer Graphics (SIGGRAPH '87 Proceedings), 21(4):163-169, July 1987
14. http://www.nersc.gov/
15. Norman M.L., Shalf J., Levy S., and Daues G. (1999) Diving deep: Data management and visualization strategies for adaptive mesh refinement simulations. Computing in Science and Engineering, 1(4):36-47, July/August 1999
16. Ousterhout J.K. (1994) Tcl and the Tk Toolkit. Addison-Wesley
17. http://www-vis.lbl.gov/
18. Schroeder W., Martin K., and Lorensen B. (1998) The Visualization Toolkit, 2nd Edition. Prentice-Hall