Astronomical Data Analysis Software and Systems VIII
ASP Conference Series, Vol. 172, 1999
D. M. Mehringer, R. L. Plante, and D. A. Roberts, eds.

Numerical Observatory: a Web Site for the Retrieval, Analysis, and Visualization of Simulated Cluster Data

Greg Daues, Christopher Currie, Peter Anninos,
Larissa Kohler, John Shalf, Michael Norman[1]

National Center for Supercomputing Applications, University of Illinois Urbana-Champaign, Urbana, IL 61801

Abstract. We have developed a user-friendly, highly interactive, "Cluster Archive" web site (`http://sca.ncsa.uiuc.edu/`) that permits users to retrieve, analyze, and visualize archived data from extremely high-resolution numerical simulations of more than 100 X-ray clusters, spanning several orders of magnitude in luminosity and mass. The Cluster Archive site provides fast and efficient client software in the form of Java applets and CGI scripts to connect with computationally intensive server software. These scripts allow fast data selection, transfer, extraction, and the execution of various analysis routines. The site also supports server side interactive 2D Java and VRML-based 3D visualizations.

1. The Digital Archive and Web Interface

The NCSA Laboratory for Computational Astrophysics, together with members of the Cosmology groups at the University of Missouri, the Massachusetts Institute of Technology, and the New Mexico State University, are performing roughly 100 extremely high resolution simulations of X-ray galaxy clusters (with several decades variance in mass). The clusters form in a self-consistent numerical treatment of dark matter and baryonic fluid with a Cold Dark Matter (CDM) initial spectrum and a flat Friedmann background cosmology (Norman & Bryan 1998). These simulations use Adaptive Mesh Refinement (AMR) to achieve unprecedented resolution in the important central cores of the clusters where the densities are highest and where X-ray emissions are the greatest. AMR works by dynamically spawning finer-grained subgrids in the regions of the domain where additional resolution is required (Berger & Oliger 1984). The dynamic and automated generation of subdomain grid hierarchies of varying sizes, resolutions, and locations complicates the data structures in terms of size and information content.

Our goal is to make available to the scientific community the results of the simulations. However, making this kind of data available presents significant technological challenges because of the size and complexity of the data structures involved. Typically the results of these simulations are thousands of multiply-embedded and cross linked AMR grid structures distributed across hundreds of

[1]Astronomy Department, UIUC

thousands of files totaling hundreds of Gbytes. Most users cannot store this amount of data on their local disk, even for a single cluster. Furthermore, they would also not have the processing speed and memory to do the analysis. In order to use the data effectively, users must have access to basic analysis and processing software on an appropriate computing server.

To meet this need, we have developed a "numerical observatory" in a workbench style system that allows users to interact with archived simulation data over the Web. This environment supports the data retrieval across several machines at the NCSA, and between those machines and the local client systems. The observatory further supports dynamic and interactive multidimensional visualization and basic analysis tools to extract specified physical attributes of the cluster systems.

When considering what type of visualization tools to provide to the client's local machine, platform independence and "thin" client visualization software (i.e., software that is small and quick to download) are critical requirements for this project. This is especially crucial for the 3D visualizations. We successfully addressed these concerns in two ways. First, we render 3D objects on the server side. Client applets downloaded from the site allow the user to specify which cluster to visualize, which field quantities to download, and the type of visualization, such as an isosurface. The applet then connects back to a CGI script via the HTTP protocol. This protocol allows the user's applet input to be POSTed to the CGI script (Kaplan 1997).

The script then takes the user's preferences and launches a server-side VTK-based rendering program. (VTK[2] is a freely available visualization toolkit for 3D computer graphics, image and volume processing, and visualization). By rendering on the server, this achieves the goal of having a "thin" client (i.e., small, fast applet) on the user's local machine.

Second, although visualizations created by VTK are not platform independent, VTK does provide classes to convert the visualizations into a VRML 2.0 format and then write them to a file. However, instead of writing to a file, our CGI scripts write the contents to the standard output and the server sends them back to the applet. The applet then simply opens a new browser window, specifies the content type (i.e., vrml), and the browser then loads the appropriate VRML-viewer plug-in. Since most browsers running under many different operating systems have the ability to display VRML-based graphical representations, this allows us to achieve a platform independent way of displaying complex 3D visualizations.

In short, the site leverages widely and freely available technology to deliver high-quality 3D visualizations to a variety of platforms without overwhelming the processing power and storage capacity of a client's local machine. In Fig. 1, we display a screenshot of a working session of the current Cluster Archive site illustrating the multiple isosurface generating features of the VRML 3D visualizations.

The requirement of lightweight client visualization software is also met by Cluster Archive site tools that implement lower-dimensional visualizations through Java graphics. Specifically, the client applet allows the user to launch

---

[2] `http://www.kitware.com/vtk.html`

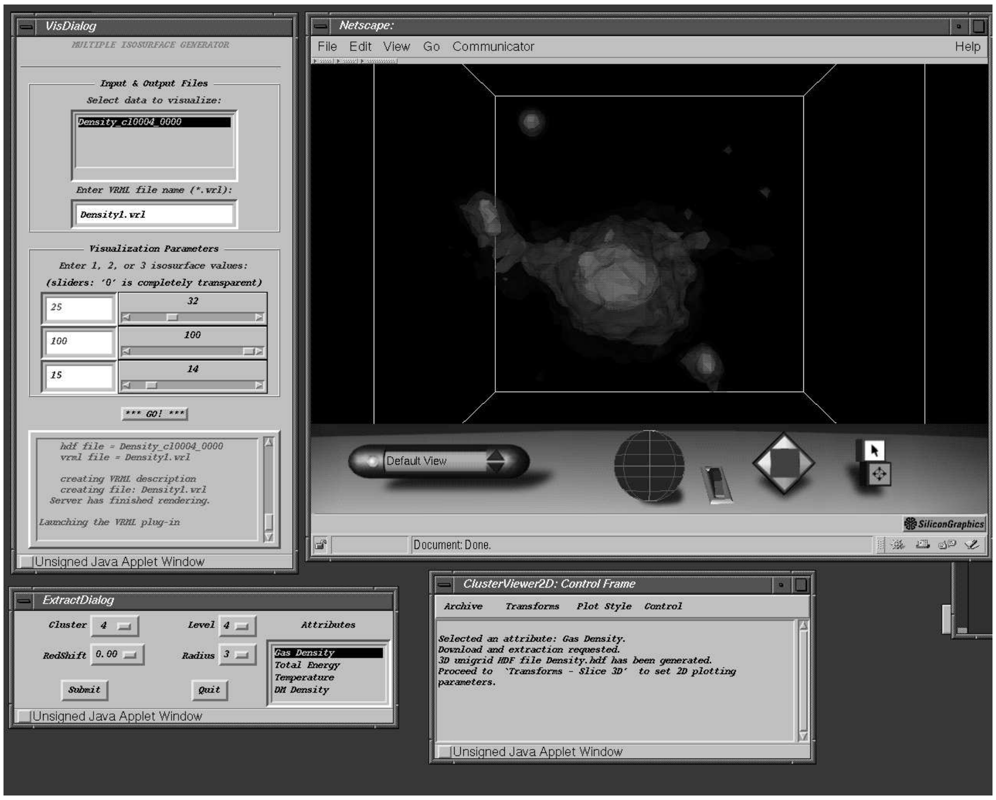

Figure 1. Screenshot of a working session of the Cluster Archive site illustrating VRML-based 3D visualization. Components of a Java applet allow user control of field quantities for data retrieval and 3D multiple isosurface visualization. Three different isosurfaces of constant gas density for a typical cluster are shown.

CGI scripts that sample subsets of large 3D Hierarchical Data Format (HDF) files that are retrieved from the archive storage and reside on the server side. Only these smaller subsets of the data are transferred over the network for processing by the client software. One capability of the client applet is the construction of 2D color contour plots of the sampled data. A screenshot displaying a working session of the web site implementing these plots is shown in Fig. 2. In keeping with the consideration of processing speed and memory on the client side, the client applet allows users to control the size and number of images generated. Archive tools for other forms of lower-dimensional visualizations such as 2D projections of 3D data and 1D line plots are also under development.

2. Future Enhancements

Future developments of the Cluster Archive site will focus on implementing Remote Method Invocation (RMI) to establish a more persistent and unified compute environment, allowing for finer controlled access of objects on the server.

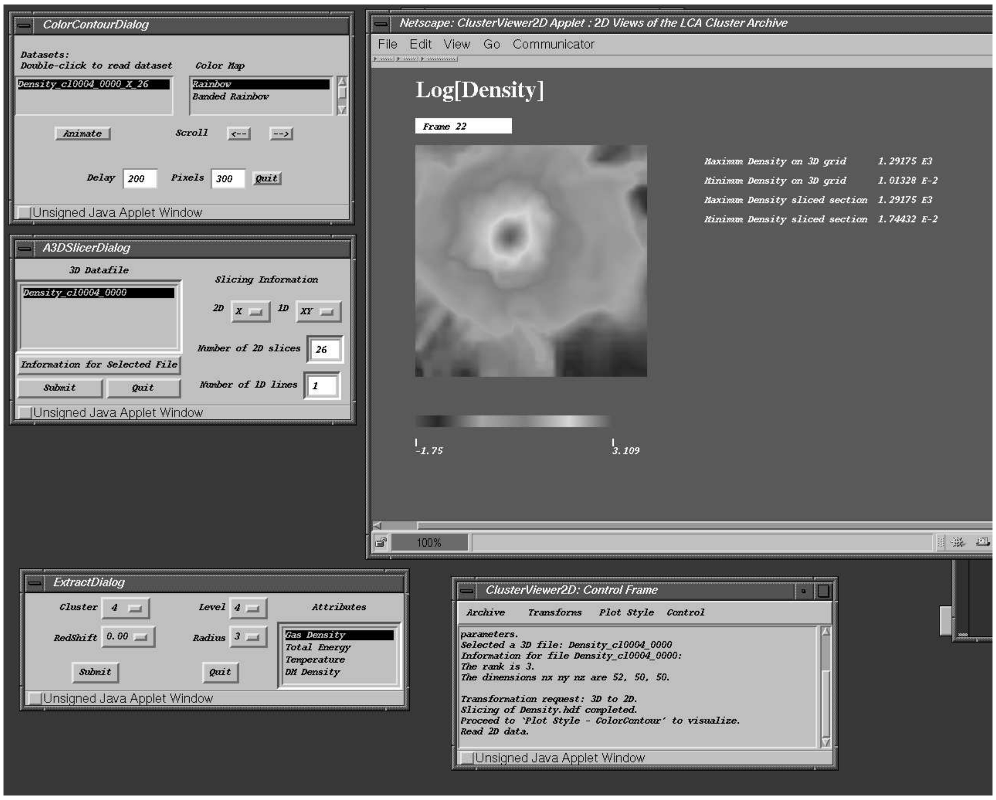

Figure 2. Screenshot of a working session of the Cluster Archive site illustrating Java-based 2D visualization. The components of the Java applet allow the user to specify parameters for data retrieval, sampling of the data into lower-dimensional subsets, and 2D color contour visualization with multiple color maps and zooming features. A plot of the gas density for a slice through the core of a typical cluster is shown.

We will also add links to other electronic literature and develop software to interact with existing digital libraries of observed data. This will necessitate adding support for more data types and formats, as well as the construction of additional analysis tools to allow direct comparison with observed data.

Acknowledgments. This work was supported in part by NASA ATP grant NAG 5-7404.

References

Berger, M. J. & Oliger, J. 1984, J. Comp. Phys. 53, 484

Kaplan, L. 1997, Java Report, 5, 51

Norman, M. & Bryan, G. 1998, Phys. Rev. Lett., 81, 3815